



PUBLIC

SAP Data Intelligence

2024-11-20

Administration Guide

Content

- 1 Administration Guide for SAP Data Intelligence Cloud. 7**
- 2 Getting Started in the Cloud. 8**
- 2.1 Create an SAP Data Intelligence Cloud Instance in SAP BTP. 8
 - Add a Tenant to an Existing SAP Data Intelligence Cloud Instance. 10
 - Enable Customer Managed Key Using SAP Data Custodian. 13
 - Enable Customer Managed Key Using JSON Parameter Visualization. 15
- 2.2 Update an SAP Data Intelligence Cloud Instance in SAP BTP. 16
- 2.3 Back Up and Restore an SAP Data Intelligence Cloud Instance. 18
- 2.4 Managing SAP Data Intelligence Cloud Cluster Hibernation and Wakeup. 18
 - Triggering Hibernation and Wakeup. 19
 - Scheduling Hibernation and Wakeup. 19
- 2.5 Configure Cloud Connector 22
- 2.6 Production Environment Maintenance. 25
- 3 Using SAP Data Intelligence Cloud System Management. 27**
- 3.1 Access SAP Data Intelligence Cloud System Management. 27
- 3.2 Managing Applications. 28
 - Launch an Application. 29
 - Start or Restart a Stable Application Instance. 30
 - Set Application Configurations. 31
 - Add an Application Secret. 31
 - SAP Data Intelligence Cloud Nested Applications. 32
 - Manage Metadata Explorer Catalog and Data Preparation Memory. 32
 - Manage Metadata Automatic Lineage Extractions. 34
 - Manage Size Limit for File Upload and Download. 36
 - Manage Metadata Rule Validation CSV Length. 37
 - Manage Metadata Failed Records. 37
 - Manage Modeler Customer Registry Secret. 43
- 3.3 Managing Users. 45
 - User Policies. 47
 - Create and Delete users and User Instances 47
 - Change a Password. 48
 - Assign Policies. 49
 - Remove Policies From a User. 50
- 3.4 Manage Policies. 50
 - Policy Resources and Resource Types 53

	Pre-Delivered Policies.	55
	Nested Policies.	58
	Application Start Policies.	60
	Resource Quotas.	64
	Create a Policy.	65
	Assign a Policy to a User	67
	Mapping Policies to Identity Providers.	67
3.5	Managing Files.	70
	Create a Folder.	71
	Create a File.	72
	Delete a File or Folder.	72
	Rename a File or Folder.	73
	Copy the Path of a File or Folder.	73
	Export a File or Folder.	74
	Import a File or Solution.	74
	Import a Solution from the Solution Repository.	75
	Export a File or Folder as a Solution.	76
	Export a Solution to Solution Repository.	76
	Sharing Files Using Solution Repository.	77
3.6	Manage Tenant.	78
3.7	Strategies.	80
3.8	System Management Application Configuration and Secrets Properties.	81
4	Using SAP Data Intelligence Connection Management.	118
4.1	Log on to SAP Data Intelligence Connection Management.	120
4.2	Create a Connection.	120
4.3	Manage Certificates.	122
4.4	Supported Connection Types.	122
	ABAP.	124
	ABAP LEGACY.	127
	ADL (Deprecated).	131
	ADL_V2.	132
	AWS_SNS.	134
	AZURE_SQL_DB.	135
	BW.	137
	CLOUD_DATA_INTEGRATION.	138
	CPEM.	140
	CPI.	141
	DATASERVICES.	142
	DB2.	143
	GCP_BIGQUERY.	145
	GCP_DATAPROC.	148

GCP_PUBSUB.	149
GCS.	149
HANA_DB.	150
HANA_XS.	153
HDFS.	154
HDL_DB.	156
HDL_FILES.	159
HTTP.	160
IMAP.	161
INFORMATION_STEWARD.	162
KAFKA.	163
MSSQL.	169
MYSQL.	170
ODATA.	171
OPEN_CONNECTORS.	173
OPENAPI.	175
ORACLE.	176
OSS.	179
POSTGRESQL.	181
REDSHIFT.	182
RSERVE.	186
S3.	186
SAP_IQ.	189
SDL.	192
SFTP.	194
SMTP.	196
SNOWFLAKE.	197
TERADATA.	201
WASB.	204
4.5 Supported Connections in SAP Data Intelligence.	205
Supported Data Source Systems.	206
Supported Data Target Systems.	215
Remote System Orchestration.	218
Replication Flow Connections.	219
4.6 Using SAP Cloud Connector Gateway.	223
4.7 (Mandatory) Configure Authorizations for Supported Connection Types.	224
HANA_DB.	224
4.8 Allowing SAP Data Intelligence Access Through Firewalls.	225
5 Monitoring SAP Data Intelligence	227
5.1 Monitoring Graphs	227
5.2 Log in to SAP Data Intelligence Monitoring.	229

5.3	Using the Monitoring Application.	229
5.4	Stop, Restart, and Pause Tenant User Graphs.	237
5.5	Accessing the SAP Data Intelligence Monitoring Query API.	238
	Access Restrictions.	240
	Metric Resolution and Retention.	240
	Retrieving Your Tenant UID.	241
	Testing the SAP Data Intelligence Monitoring Query API.	242
	Running PromQL Instant Queries.	243
	Running PromQL Range Queries.	244
	Using a Series Request to List Available Metrics	246
	Accessing the SAP Data Intelligence Monitoring Query API Via POST Requests.	247
	Advanced Query Expressions.	248
6	Integration Monitoring in SAP Cloud Application Lifecycle Management.	252
6.1	Using the SAP Cloud ALM Integration Monitoring for SAP Data Intelligence.	255
7	Maintaining SAP Data Intelligence.	257
7.1	On-Demand Certificate Renewal.	257
	Functional Cluster.	258
	Non-Functional Cluster.	258
8	Exporting Customer Data.	259
8.1	Log in to SAP Data Intelligence Customer Data Export	261
9	Improving Performance.	262
9.1	Improving CDC Graph Generator Operator Performance.	262
10	Sizing for Metadata Explorer and Self-Service Data Preparation.	264
10.1	Configure App-Data.	264
10.2	Configure Other Applications.	266
11	Understanding Security.	267
11.1	Data Protection and Privacy in SAP Data Intelligence.	267
	Managing Audit Logs.	270
	Viewing Audit Logs.	271
	Malware Scanning.	272
11.2	Security Recommendations for SAP Data Intelligence.	272
11.3	Securing SAP Data Intelligence.	277
	Enabling Authentication for SAP Data Intelligence Services and Users.	278
	Configuring External Identity Providers in SAP Data Intelligence.	278
	Giving User Permissions for SAP Data Intelligence Access.	279
	Connecting Your On-Premise Systems To SAP Data Intelligence.	280
12	Troubleshooting SAP Data Intelligence.	284

12.1	Troubleshooting SAP Cloud Connector.	284
12.2	Troubleshooting Flowagent.	285

1 Administration Guide for SAP Data Intelligence Cloud

This Administration Guide contains information for administrators about configuring, monitoring, and managing SAP Data Intelligence Cloud.

Related Information

[Getting Started in the Cloud \[page 8\]](#)

[Using SAP Data Intelligence Cloud System Management \[page 27\]](#)

[Using SAP Data Intelligence Connection Management \[page 118\]](#)

[Manage Policies \[page 50\]](#)

[Monitoring Graphs \[page 227\]](#)

Using SAP Data Intelligence Diagnostics

Using SAP Data Intelligence Metrics

[Maintaining SAP Data Intelligence \[page 257\]](#)

[Exporting Customer Data \[page 259\]](#)

[Improving Performance \[page 262\]](#)

[Understanding Security \[page 267\]](#)

[Troubleshooting SAP Data Intelligence \[page 284\]](#)

2 Getting Started in the Cloud

Before you and your tenants start to use SAP Data Intelligence Cloud, you must create and configure basic administrative features, such as instances, clusters, and the cloud connector.

Related Information

[Create an SAP Data Intelligence Cloud Instance in SAP BTP \[page 8\]](#)

[Update an SAP Data Intelligence Cloud Instance in SAP BTP \[page 16\]](#)

[Back Up and Restore an SAP Data Intelligence Cloud Instance \[page 18\]](#)

[Managing SAP Data Intelligence Cloud Cluster Hibernation and Wakeup \[page 18\]](#)

[Configure Cloud Connector \[page 22\]](#)

[Production Environment Maintenance \[page 25\]](#)

2.1 Create an SAP Data Intelligence Cloud Instance in SAP BTP

Before you and your tenants can access and use SAP Data Intelligence Cloud, you must create a new instance in the SAP Business Technology Platform (BTP) cockpit.

Prerequisites

Before you create an instance in the SAP BTP cockpit, ensure the following:

- Your global account has a commercial entitlement with either cloud credits (consumption-based model) or a subscription contract.
- You're assigned the Space Developer role in the BTP space by your sub account administrator or your space manager.
- You use Google Chrome so that you can view popups properly in SAP BTP.

Context

In this guide, we provide high-level steps to create an SAP Data Intelligence Cloud instance on SAP BTP. For more detailed information, or for instructions that use the SAP BTP Cloud Foundry command line interface, see the [SAP BTP documentation](#).

Procedure

1. Create a subaccount in your global account by performing the following substeps:
 - a. Open your SAP BTP cockpit, and choose *Global Accounts*.
 - b. Select the global account that is entitled for SAP Data Intelligence Cloud.
 - c. Choose *New Subaccount*.
 - d. Enter or select the subaccount information.

The SAP Data Intelligence Cloud subaccount must contain the following values:

Option	Value
Environment	Cloud Foundry
Provider	AWS
Region	Select the applicable region: <ul style="list-style-type: none">• <i>AWS</i>: Australia (Sydney), Europe (Frankfurt), Japan (Tokyo), Singapore, US East (VA)• <i>Microsoft Azure</i>: Europe (Netherlands), US West (WA)
Subdomain	Name that becomes part of the URL. The subdomain name can contain any of the following characters: <ul style="list-style-type: none">• Letters• Digits• Hyphens (but not as the first and last characters) The subdomain name must be unique across all accounts in the same region of the SAP BTP Cloud Foundry environment.

- e. Choose *Create*.

A new Cloud Foundry subaccount tile appears on the Global Accounts page with the subaccount details.

2. Choose *Enable Cloud Foundry*, and return to the Global Accounts overview page.
3. Choose **► Entitlements ► Subaccount Assignment ►**, and then choose **► Configure Entitlements ► Add Service Plans ►**.
4. Add the SAP Data Intelligence Cloud entitlement for the subaccount.

The subaccount lists the available quota for SAP Data Intelligence Cloud.
5. Create a space in your subaccount by performing the following substeps:
 - a. Choose **► Subaccount Spaces ► New Space ►**
 - b. Enter the space name.
 - c. Select the permissions to assign to your ID.

A space carries the quota and consumes all available quota provided by the cloud credits.
6. Create an SAP Data Intelligence Cloud instance by performing the following substeps:
 - a. Choose **► Services ► Service Marketplace ►**.

- b. Search for “SAP Data Intelligence Cloud”, and select the SAP Data Intelligence Cloud service.

Note

A subaccount supports only one SAP Data Intelligence Cloud instance.

- c. Choose ► [Instances](#) ► [New Instance](#). ►
- d. Follow the application's steps to create an instance, specifying the desired credentials and sizing configuration.
- e. Choose the instance type of subscription (dedicated).

Restriction

CPEA (enterprise) is a deprecated service plan. We recommend that you switch to the dedicated service plan when your service is up for renewal.

- f. Choose a minimum and maximum number of worker nodes. SAP Data Intelligence Cloud scales based on the usage.

To use VPN or VPC peering connectivity when your network has an IP address conflict with 10.0.0.0/16, specify a CIDR (Classless Inter-Domain Routing) block value (/22 or larger) for your SAP Data Intelligence Cloud network.

Note

The SAP Data Intelligence Cloud network uses the default CIDR value 10.0.0.0/16.

SAP BTP creates the instance, however instance creation can take up to an hour.

7. When the instance status states “Created”, choose ► [Actions](#) ► [Open Dashboard](#). ►
8. Log on to SAP Data Intelligence Cloud on the default tenant with the credentials that you provided during instance creation.

Related Information

[Add a Tenant to an Existing SAP Data Intelligence Cloud Instance \[page 10\]](#)

[Enable Customer Managed Key Using SAP Data Custodian \[page 13\]](#)

[Enable Customer Managed Key Using JSON Parameter Visualization \[page 15\]](#)

[Connect Using Site-to-Site VPN \[page 280\]](#)

[Connect Using Virtual Network Peering \[page 282\]](#)

2.1.1 Add a Tenant to an Existing SAP Data Intelligence Cloud Instance

If you have a running SAP Data Intelligence Cloud instance, you can add new logically isolated tenants to it.

Prerequisites

Before you perform this task, ensure that your global account has a commercial entitlement through either cloud credits (consumption-based model) or a subscription contract. Also ensure that you have a running SAP Data Intelligence Cloud instance.

After you create an instance, choose a service plan to create an instance of the SAP Data Intelligence Cloud service. SAP Data Intelligence Cloud provides the service plans described in the following table.

Service Plan	Description
Dedicated	<p>An isolated SAP Data Intelligence Cloud cluster with dedicated hardware resources, and with one tenant named "default", which the application creates automatically.</p> <p>When you configure the sizing for your Dedicated service plan, keep in mind that additional tenants share the same hardware resources of their parent instance. A heavy workload on one tenant can slow workloads running on another tenant.</p> <p>To limit how many resources a tenant can use, specify resource quotas when you create the tenant.</p>
Tenant	<p>The application uses the Tenant service plan in combination with the Dedicated service plan's "default" tenant. In addition to the default tenant, the tenant plan allows you to add up to 19 tenants to a given SAP Data Intelligence Cloud instance. The tenants run on the same cluster and share the same hardware resources, but are isolated from each other.</p> <p>Add new isolated tenants to form logical partitions inside an SAP Data Intelligence Cloud instance. Data from one tenant isn't accessible to other tenants. Each tenant has different users and data. For example, configure the plan to use one isolated tenant per department of your business.</p>

Note

The service plans vary based on your account type. If you are still using a hybrid-based account, where you have access to subscription and CPEA (enterprise) cloud credits, use the Dedicated plan to avoid using your CPEA cloud credits for the added tenant.

Restriction

The CPEA (enterprise) plan is deprecated. Therefore, SAP recommends switching your plan to the dedicated service plan when your service is up for renewal.

Context

To add new tenants to your existing SAP Data Intelligence Cloud instance, perform the following steps:

Procedure

1. Open the SAP BTP cockpit and choose [Global Accounts](#).
2. Select the global account, and then the subaccount, that contains your SAP Data Intelligence Cloud instance.
3. Select the space where you created your SAP Data Intelligence Cloud enterprise instance from the spaces list.
4. Choose [► Services ► Service Marketplace ►](#).
5. Search for “Data Intelligence” and select the SAP Data Intelligence Cloud service.
6. Choose the *tenant* plan in [Choose Service Plan](#), and choose [Next](#).

The [Specify Parameters](#) dialog box opens.

7. Select the name of your tenant, and provide the associated user name and password.
8. **Optional:** Add resource quotas to limit how many resources the tenant can use.

You can have multiple quotas, and specify whether the target to be limited belongs to the application or the workload. Apply quotas to the following three types of resources:

- CPU (millicpu)
- Memory (megabytes)
- Pod count (number of pods)

9. If you've more than one SAP Data Intelligence Cloud dedicated service plan instance, select the applicable instance from the [Cluster Name](#) list and choose [Next](#).

If you have only one dedicated service plan instance, the system selects the one instance automatically in [Cluster Name](#).

10. Choose your instance name and select [Finish](#).

The creation of a new tenant can take up to 20 minutes.

To use the new tenant, return to the instance list view.

Next Steps

After the tenant [Last Operation](#) field displays “Created”, choose [Open Dashboard](#) in the Actions column. The system redirects you to the tenant login screen. SAP automatically completes the tenant name and user credentials that you selected when you added the tenant, so that you can log on to the new tenant.

2.1.2 Enable Customer Managed Key Using SAP Data Custodian

For SAP Data Intelligence Cloud landscapes hosted in Amazon Web Service (AWS), use the Customer Managed Key (CMK) feature to integrate an SAP Data Custodian key when you create an SAP Data Intelligence Cloud instance in the SAP Business Technology Platform (BTP).

Context

The following table describes the implementation methods for integrating an SAP Data Custodian key when you create an SAP Data Intelligence Cloud instance in SAP BTP.

Implementation Method	Description
CSEK (Customer Specific Encryption Keys)	Use SAP Data Custodian as an encryption tool in the SAP Data Intelligence Cloud instance, but don't provide an existing SAP Data Custodian Key ID. SAP Data Intelligence Cloud generates a new SAP Data Custodian Key in the SAP Data Custodian tenant.
CCEK (Customer Controlled Encryption Keys)	You provide a specific non-HYOK (Hold Your Own Key) that you own to use for encryption purposes in the new SAP Data Intelligence Cloud instance. You can use the same key multiple times for different SAP Data Intelligence Cloud instances.
HYOK (Hold Your Own Key)	You provide an SAP Data Custodian HYOK key that you own to use for encryption purposes in the new SAP Data Intelligence Cloud instance. For HYOK, you must perform some additional steps to modify your AWS Key Management System Key policy so that SAP Data Intelligence Cloud can use the HYOK.

Restriction

The following procedure applies only to SAP Data Intelligence Cloud landscapes hosted in AWS. SAP Data Intelligence Cloud instances hosted in other cloud providers don't have access to the CMK feature.

To create SAP Data Intelligence Cloud instances that use the CMK feature, perform the following steps:

Procedure

1. Obtain your SAP Data Custodian Key ID.

The SAP Data Custodian Key ID isn't the same ID as the key's DKR (Dynamic Key References).

2. Start to create a new SAP Data Intelligence Cloud instance through SAP BTP.

When *Use CMK* is enabled, the *Key Reference* field becomes visible.

- Complete the *Key Reference* field as instructed in the following table based on your implementation method.

Implementation Method	Step
CSEK	<p>Leave the <i>Key Reference</i> field empty.</p> <p>This is the final step in the process. Skip the remaining steps.</p>
CCEK	<p>Pass an SAP Data Custodian Key ID to the <i>Key Reference</i> field.</p> <p>This is the final step in the process. Skip the remaining steps.</p>
HYOK	<p>Pass an SAP Data Custodian Key ID to the <i>Key Reference</i> field, and complete the remaining steps.</p>

- For HYOK method only.** Open the SAP Data Custodian user interface and find your HYOK key.

It takes a few minutes for the system to create the key, then the system adds a new label in the named SAP Account.

- For HYOK method only.** Request the AWS account ARN (Amazon Resource Name) information for the new SAP Data Intelligence Cloud instance:
 - Create an SAP support ticket using the component **CA-DI-SI**.
 - After you receive the ARN from support, use it in the next step for the key policy.

The ARN goes in `<SAP_ACCOUNT_LABEL_VALUE>`.

Sample Code

```
"AWS": [
  "<SAP_ACCOUNT_LABEL_VALUE>"
]
```

- For HYOK method only.** Complete the key policy by performing the following substeps:
 - Open the *AWS KMS Key* panel in the AWS console.
 - Select the provided key in the *Customer managed keys* category to open its properties.
 - Choose *Switch to policy view* in the *Key policy* category and choose *Edit*.
 - In the key policy, add the following code entry, using the ARN that you receive from SAP support:

```
{
  "Sid": "Allow SAP DI to use this KMS Key",
  "Effect": "Allow",
  "Principal": {
    "AWS": [
      "<SAP_ACCOUNT_LABEL_VALUE>"
    ]
  },
  "Action": [
    "kms:CreateGrant",
    "kms:ListGrants",
    "kms:RetireGrant",
    "kms:RevokeGrant",
    "kms:DescribeKey",
    "kms:GetKeyRotationStatus",
    "kms:GetPublicKey",
  ]
}
```

```
"kms:Encrypt",
"kms:Decrypt",
"kms:ReEncrypt",
"kms:Sign",
"kms:Verify",
"kms:GenerateRandom",
"kms:GenerateDataKey*"
],
"Resource": "*"
}
```

Note

There can be only one instance of the policy for each of the SAP accounts that use the key. Therefore, if the key has already been used for SAP Data Intelligence Cloud encryption, it's possible that the policy already exists for the specific AWS SAP account. If the key has already been used, don't add the policy a second time.

Results

Your instance creates and uses the key provided for encrypting SAP Data Intelligence Cloud resources in AWS.

2.1.3 Enable Customer Managed Key Using JSON Parameter Visualization

Enable the Use Customer Managed Key (CMK) value and populate the Key Reference parameter of your instance using the JSON parameter visualization method.

Context

The steps in [Enable Customer Managed Key Using SAP Data Custodian \[page 13\]](#) use the Form parameter visualization in SAP Business Technology Platform (BTP). In the following steps, to activate the Use CMK value and populate the Key Reference parameter of your instance, you use the JSON Parameter Visualization.

Procedure

1. Obtain your SAP Data Custodian Key ID.

Don't confuse the SAP Data Custodian Key ID with the key's DKR (Dynamic Key References).

2. Start creating a new SAP Data Intelligence instance through SAP BTP.
3. Enable the *Use CMK* field by adding the following line to the end of the JSON parameter list:

```
"useCMK": "true"
```

4. To use the *Key Reference* field, enter the following line after the "useCMK" line:

```
"keyRef": "<YOUR_KEY_REFERENCE>"
```

Enter a *Key Reference* value for the CCEK (Customer Controlled Encryption Keys) and HYOK (Hold Your Own Key) implementation methods.

For the Customer Specific Encryption Keys (CSEK) method, leave the *Key Reference* field blank, and enter the following line:

```
"keyRef": ""
```

Results

If you enable CMK through the JSON parameter list as instructed, your parameters list in JSON looks similar to the following example:

Sample Code

```
{
  "minNodes": "1",
  "maxNodes": "1",
  "adminUsername": "",
  "adminPassword": "",
  "vpcCIDR": "10.0.0.0/16",
  "hibernationSchedules": "[]",
  "useCMK": "true",
  "keyRef": "<YOUR_KEY_REFERENCE>"
}
```

2.2 Update an SAP Data Intelligence Cloud Instance in SAP BTP

To extend the capacity or reduce costs of your SAP Data Intelligence Cloud instance, update the running instance in SAP Business Technology Platform (BTP).

Prerequisites

Before you perform the following steps, ensure that you have a running SAP Data Intelligence Cloud instance in SAP BTP.

Context

This guide contains high-level steps to update an SAP Data Intelligence Cloud instance in SAP BTP. For more detailed information, or for instructions that use the Cloud Foundry command line Interface, see the [SAP BTP documentation](#).

To update a running SAP Data Intelligence Cloud instance in SAP BTP, perform the following steps:

Procedure

1. Open the SAP BTP dashboard and choose the subaccount where the SAP Data Intelligence Cloud instance was created.
2. Open *Instances and Subscriptions* and find the SAP Data Intelligence Cloud instance with the service "SAP Data Intelligence".
3. Choose **Actions > Update**.
4. **Optional:** Update the instance service plan in the *Plan* list, and follow the provided steps.

Update service plans as follows:

- From `<enterprise>` to `<dedicated>`.
- From `<dedicated>` to `<enterprise>`. (The Enterprise plan is deprecated, so SAP doesn't recommend this update type.)

Note

It's not possible to update a `<tenant>` service plan to `<enterprise>` or `<dedicated>`, nor can you update an `<enterprise>` or `<dedicated>` service plan to `<tenant>`.

5. Open the *Parameters* tab.
6. Specify the applicable sizing configuration and hibernation configuration as described in the following table.

Sizing Configuration	Hibernation Configuration
You can update the <i>Min Number of Worker Nodes</i> and <i>Max Number of Worker Nodes</i> by updating an SAP Data Intelligence Cloud instance.	To adjust hibernation, see the following topics:
Decreasing the <i>Max Number of Worker Nodes</i> triggers a restart of the instance.	<ul style="list-style-type: none">• Managing SAP Data Intelligence Cloud Cluster Hibernation and Wakeup [page 18]• Scheduling Hibernation and Wakeup [page 19]

7. Choose *Update Instance*.

When the instance status is green, your instance is updated.

Related Information

[Managing SAP Data Intelligence Cloud Cluster Hibernation and Wakeup \[page 18\]](#)

[Scheduling Hibernation and Wakeup \[page 19\]](#)

2.3 Back Up and Restore an SAP Data Intelligence Cloud Instance

SAP Data Intelligence Cloud service supports backup and restore concepts.

For information about backing up and restoring an SAP Data Intelligence Cloud instance, see SAP Note [3153355](#).

2.4 Managing SAP Data Intelligence Cloud Cluster Hibernation and Wakeup

To reduce costs while your SAP Data Intelligence Cloud clusters aren't in use, you can "pause" cluster nodes and trigger hibernation.

During hibernation, SAP Data Intelligence Cloud stops all cluster nodes and persists only storage. Hibernation provides many advantages from a cost perspective, and costs are highly reduced, but not eliminated.

Note

To trigger hibernation and add schedules, you must be in the SAP Business Technology Platform (BTP) Cloud Foundry space where you created SAP Data Intelligence Cloud. You must also have either a Space Manager or Space Developer role. For information about adding space users, see [Add Space Members Using the Cockpit](#) in the SAP BTP documentation.

The hibernation process takes approximately 15 minutes to complete; the wake-up process takes approximately 30 minutes.

SAP Data Intelligence Cloud doesn't trigger hibernation under the following circumstances:

- While graphs are running. (You must first manually stop them).
- While an operational task, such as backup, is in progress.

You can't run the following processes during hibernation:

- Update or delete clusters.
- Backups.
- Scheduled Modeler jobs.

Note the following additional information about hibernation:

- During hibernation, replication flows that replicate changes from remote system objects, such as ABAP and DB tables, are stopped. If this type of replication isn't stopped, it can cause continuous growth of change data in the remote systems, which can impact remote system performance and availability.
- During hibernation, the cluster URL is not accessible, and the system issues an error that the connection timed out.
- During a maintenance window, the system wakes up hibernating clusters and puts them back into hibernation only after maintenance completes.
- Availability tracking doesn't consider the hibernation period.
- When a cluster is in hibernation, choosing [View Dashboard](#) redirects the application to an informational page.

2.4.1 Triggering Hibernation and Wakeup

To trigger hibernation and wakeup in SAP Data Intelligence, use the command line.

Note

You can also trigger and wakeup hibernation in the user interface, however, these instructions are just for command line.

To trigger hibernation, use a command line like the following example:

Sample Code

```
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"enableHibernation":"true"}'
```

To wakeup from hibernation, use a command line like the following example:

Sample Code

```
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"enableHibernation":"false"}'
```

2.4.2 Scheduling Hibernation and Wakeup

In addition to manually triggering hibernation, you can also schedule hibernation and cluster wakeup by configuring hibernation schedules.

Note

SAP Data Intelligence Cloud triggers scheduled hibernation only when graphs or operational tasks are running.

Configure hibernation and wake up scheduling using one of the following methods:

- cf (Cloud Foundry) command line
- SAP Business Technology Platform (BTP) user interface

You can configure up to seven schedules simultaneously. When you use the SAP BTP user interface, configure schedules using the browser's time zone. When you use the cf command line, you must provide schedules in UTC.

Configure the hibernation and wake up schedule with the command line using cron expressions.

Note

cron doesn't allow the fields Day of Month and Month.

In the cf command line, use the properties start and end to configure hibernation and wake up respectively.

Example

The following commands show an example of hibernation and wake up schedule configuration in cf command line:

```
# Configuring hibernation schedule from monday to friday at 18h UTC
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"hibernationSchedules": "[{"start": "\0 18 * * 1,2,3,4,5"}]"}'
```

```
# Configuring wake up schedule from monday to friday at 8h UTC
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"hibernationSchedules": [{"end": "\0 8 * * 1,2,3,4,5"}]"}'
```

```
# Configuring hibernation and wake up schedules
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"hibernationSchedules": [{"start": "\0 18 * * 1,2,3,4,5", "end": "\0
8 * * 1,2,3,4,5"}]"}'
```

```
# Configuring hibernation at 18h UTC on Mondays and wake up at 8h UTC on
Tuesdays
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"hibernationSchedules": [{"start": "\0 18 * * 1"}, {"end": "\0 8 * *
2"}]"}'
```

```
# Invalid as Day of Month and Month are not allowed
cf update-service $DATA_INTELLIGENCE_INSTANCE_NAME -p enterprise -c
'{"hibernationSchedules": [{"end": "\0 8 * 1 1"}]"}'
```

The following examples show additional scheduling configured in the user interface.

Example

Hibernate at 00:00 every Saturday and wake up at 00:00 every Monday (hibernate from 00:00 Saturday to 00:00 Monday; that is, the whole weekend):

Configure Hibernation Schedule



<input type="checkbox"/>	Weekdays	Hibernate at	Wake up at
<input type="checkbox"/>	Sat <input type="button" value="x"/> <input type="button" value="v"/>	00:00 <input type="button" value="⌚"/>	HH:mm <input type="button" value="⌚"/>
<input type="checkbox"/>	Mon <input type="button" value="x"/> <input type="button" value="v"/>	HH:mm <input type="button" value="⌚"/>	00:00 <input type="button" value="⌚"/>

☼ Example

Hibernate at 20:00 every day and wake up at 06:00 every day (hibernate from 20:00 to 06:00 every day; that is, every night):

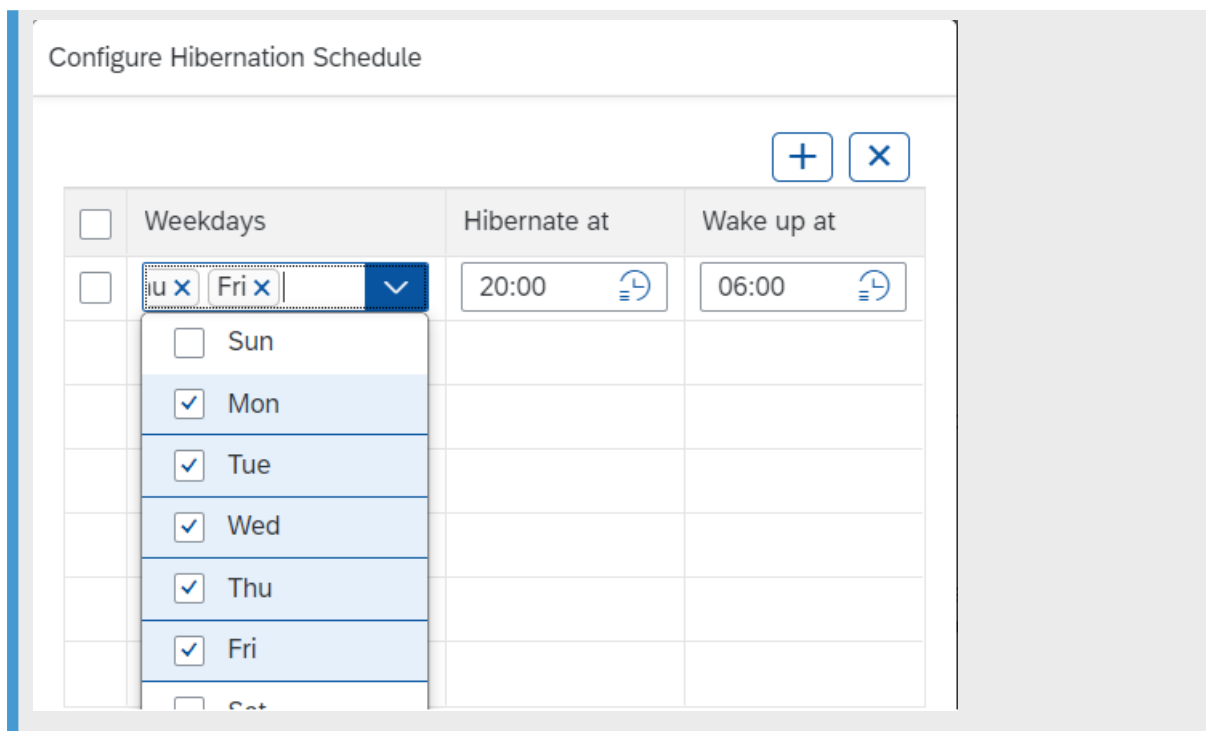
Configure Hibernation Schedule



<input type="checkbox"/>	Weekdays	Hibernate at	Wake up at
<input type="checkbox"/>	<input type="button" value="v"/>	20:00 <input type="button" value="⌚"/>	06:00 <input type="button" value="⌚"/>

☼ Example

Hibernate at 20:00 and wake up at 06:00 every Monday to Friday (hibernate from 20:00 Friday to 06:00 Monday and from 20:00 to 06:00 Monday to Friday; that is, every night during workdays, and the whole weekend):



2.5 Configure Cloud Connector

Use the cloud connector to establish a set of on-premise systems to make available to SAP Data Intelligence Cloud.

Prerequisites

📌 Note

Cloud connector is a part of the SAP Connectivity service that runs on SAP Business Technology Platform (BTP).

Before you configure the cloud connector, complete the following tasks:

- Install cloud connector. For more information about installing cloud connector, see [Installation](#) in the *SAP BTP Connectivity* guide. Also, for version information, see SAP Note [3302250](#) “Cloud Connector Support Strategy”.
- Install SAP Data Intelligence Cloud, which provides one entitlement for the connectivity proxy service. The proxy service is installed automatically as a service instance next to SAP Data Intelligence Cloud.
- Establish the cloud connector connection. This task must be performed by the subaccount user from the initial SAP BTP setup.

Context

The cloud connector serves as a link between SAP BTP applications, such as SAP Data Intelligence Cloud, and on-premise systems. For more information about configuring and using the cloud connector, see [Configuration](#) in the *SAP BTP Connectivity* guide.

Procedure

1. Log on to the cloud connector.
2. Choose [Add Subaccount](#) and enter or select information as described in the following table.

Field	Description
Region	Choose the region where the SAP Data Intelligence Cloud service is deployed.
Subaccount	Enter the values obtained when you registered your account on SAP BTP. To get your subaccount ID, see Find Your Subaccount ID (Cloud Foundry Environment) in the <i>SAP BTP Connectivity</i> guide.
Subaccount User and Password	User name and password of the user who approved cloud connector access to SAP BTP. The specified user must be an administrator of the subaccount.
Location ID	The cloud connector over which the connection is opened.

3. Choose [Cloud to On-Premise](#) in the main menu.

The **Cloud to On-Premise** dashboard displays the SAP Data Intelligence Cloud on-premise connections to be visible within the cloud connector network.

The internal host specifies the host and port under which SAP Data Intelligence Cloud can reach the target system within the intranet. The connection must meet the following criteria:

- Must be an existing network address that can be resolved on the intranet.
- Must have network visibility for the cloud connector.
- Must be a real address because the cloud connector tries to forward the request to the network address specified by the internal host and port.

The virtual host is the name that is displayed in SAP Data Intelligence Cloud. The fields are prepopulated with the values of the internal host and internal port. You can assign a different port to the virtual host than the internal host.

4. Establish the set of systems to be made available by the cloud connector by performing the following substeps:
 - a. Create a connection by choosing [Add](#).
 - b. Select a system from [Back-End Type](#) and choose [Next](#).
 - c. Select an SAP Data Intelligence Cloud protocol from [Protocol](#) as shown in the following table.

Supported Back-End Type	Supported Protocol
ABAP System	RFC
SAP HANA	TCP
SAP IQ	
Non-ABAP System	HTTP
	HTTPS

Note

For establishing end-to-end encryption between SAP Data Intelligence Cloud and the target system (either SAP HANA or SAP IQ), first select *Use TLS* in the connection definition and then provide the necessary server certificates in SAP Data Intelligence Cloud Connection Management. For information about certificates, see [Manage Certificates \[page 122\]](#).

Some protocols require additional configuration (for example, HTTP). For additional configuration steps, see [Initial Configuration](#) in the *SAP BTP Connectivity* guide.

Additional Information

- Consider the following information when the target requires TLS using self-signed certificates:
 - For HTTP and HTTPS connections, certificates must be stored in the cloud connector.
 - For connections tunneled with socks5 through the cloud connector (SAP HANA), certificates must be stored in SAP Data Intelligence Cloud Connection Management.
- SAP Data Intelligence Cloud does not support the Trusted Applications feature of the cloud connector.

Results

After you configure your connections in the cloud connector, complete the configuration in SAP Data Intelligence Cloud Connection Management.

Related Information

[Using SAP Cloud Connector Gateway \[page 223\]](#)

[SAP Cloud Connector](#)

[Troubleshooting SAP Cloud Connector \[page 284\]](#)

2.6 Production Environment Maintenance

To provide new features and fixes, SAP performs regular planned updates to the SAP Data Intelligence Cloud service production environment. During the update, your system is unavailable.

SAP plans maintenance during times intended to minimize your inconvenience. You can subscribe to maintenance notifications at the [Cloud Availability Center](#).

Read about the process in the following sections so that you know what to expect before, during, and after the maintenance window.

Before the Update

At the start of the maintenance window, SAP does the following:

- Pauses graphs that run with state management (Generation 2 graphs with snapshots enabled) or have auto restart configured.
- Stops all other graphs and scheduled workloads.
- Wakes up clusters that are in hibernation.

Note

When a source system is unreachable due to network issues or down due to maintenance, replication tasks executed at that time might fail after exhausting retries. If this happens, you must manually resume the replication tasks.

To avoid this situation, we recommend that you suspend any running Replication flows before doing maintenance and then resume them once maintenance is done.

During the Update

The SAP Data Intelligence Cloud service is not available for normal usage during the update.

After the Update

After the update, SAP does the following to your graphs and clusters:

- Resumes running graphs that are paused because of the update.

Note

Graphs that were already in a paused state before the upgrade started (because of user action) are not affected.

- Does **not** restart stopped graphs after the update is complete. You must manually restart graphs as applicable.
- Resumes scheduled graphs based on the schedule.
- Places all clusters that it woke from hibernation before the update back into hibernation after a successful update.

Clear Your Browser Cache

After your SAP Data Intelligence Cloud service product environment is updated successfully, SAP recommends that you clear the browser cache and access the SAP Data Intelligence Cloud user interface for a better experience.

For instructions for clearing the cache, see the documentation for your browser.

3 Using SAP Data Intelligence Cloud System Management

The SAP Data Intelligence Cloud System Management application allows you to manage applications, users, and files. It provides the initial point of access to the user-facing applications running on its server.

The following list contains some of the tasks that you can perform in the System Management application:

- Manage tenant administrator and member users within the tenant
- Manage user secrets
- Create application instances (deprecated)
- Delete application instances
- Launch applications
- Configure applications
- Manage files

Related Information

[Access SAP Data Intelligence Cloud System Management \[page 27\]](#)

[Managing Applications \[page 28\]](#)

[Managing Users \[page 45\]](#)

[Manage Policies \[page 50\]](#)

[Managing Files \[page 70\]](#)

[Manage Tenant \[page 78\]](#)

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

[System Management Application Configuration and Secrets Properties \[page 81\]](#)

3.1 Access SAP Data Intelligence Cloud System Management

Access the System Management application after you log on to SAP Data Intelligence Cloud.

Procedure

1. Enter the SAP Data Intelligence Cloud URL in your browser.

2. Enter your login credentials for SAP Data Intelligence Cloud.

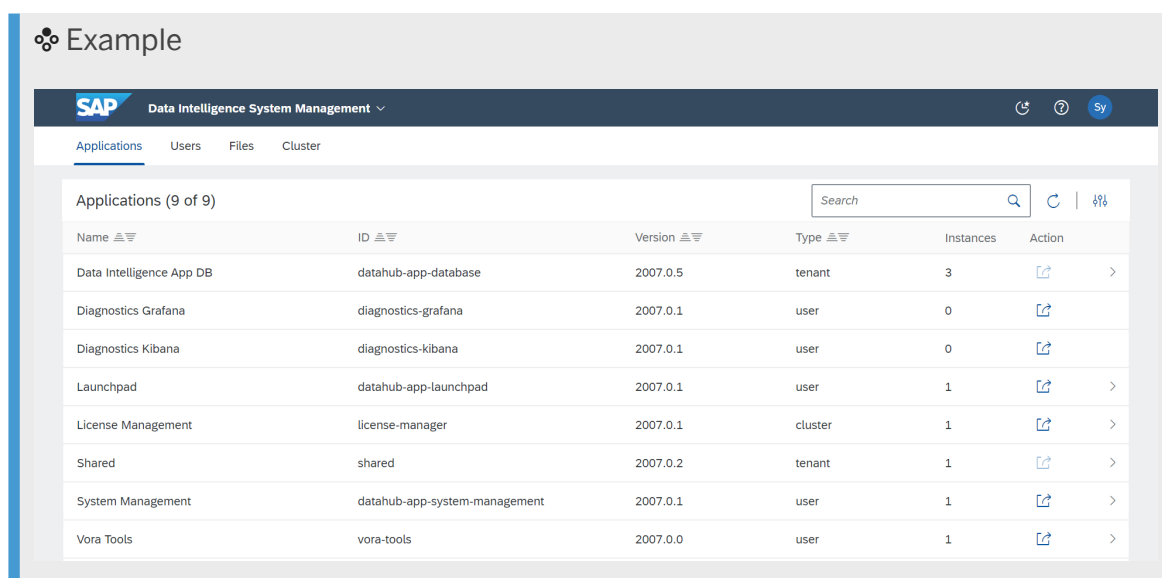
The SAP Data Intelligence Cloud Launchpad home page opens, displaying the applications as tiles. The available applications are based on the policies assigned to you.

Note

If you enter an incorrect password five consecutive times within a minute, SAP locks your account temporarily. Wait for 10 seconds before your next attempt.

3. Select the *System Management* tile.
The System Management page opens.

The following image shows an example of the initial System Management screen for member users (regular tenants). Tenant administrators can manage applications, member users, and files using the tabs at the top of the page. Member users can manage their own user information and files.



3.2 Managing Applications

The application management feature of SAP Data Intelligence Cloud System Management allows you to launch stable instances of applications associated with tenants. For example, you can launch instances of Launchpad, Modeler, and Connection Management.

Both tenant administrators and member users can manage applications. The management capabilities are based on your user type. For more information about user types, see [Managing Users \[page 45\]](#).

The System Management *Applications* page lists all applications currently available for the active tenant. When you're in the *Application* page, you can take the following actions:

- Start or restart an application instance based on your assigned permissions.
- Access an instance of an application launched from application management.
- Start or restart a stable application instance.

- [Configure applications \(tenant administrators only\)](#).

Related Information

[Launch an Application \[page 29\]](#)

[Start or Restart a Stable Application Instance \[page 30\]](#)

[Set Application Configurations \[page 31\]](#)

[Add an Application Secret \[page 31\]](#)

[SAP Data Intelligence Cloud Nested Applications \[page 32\]](#)

[Manage Metadata Explorer Catalog and Data Preparation Memory \[page 32\]](#)

[Manage Metadata Automatic Lineage Extractions \[page 34\]](#)

[Manage Size Limit for File Upload and Download \[page 36\]](#)

[Manage Metadata Rule Validation CSV Length \[page 37\]](#)

[Manage Metadata Failed Records \[page 37\]](#)

[Manage Modeler Customer Registry Secret \[page 43\]](#)

3.2.1 Launch an Application

Launch applications from the [Applications](#) page in SAP Data Intelligence Cloud System Management.

Prerequisites

If the application doesn't have a launch option, it can mean that the application doesn't have a user interface. Applications can be launched only when they have a user interface.

Context


When you open the [Applications](#) page, all applicable applications appear in rows on the page. At the end of each row, the [Action](#) column contains action icons.

Procedure

1. Locate the row for the applicable application.

→ Tip

Use the search feature if you can't find the application by scanning the list.

2. Choose  (*Launch Application*) in the *Action* column for the application.

Results

The applicable application opens.

Example

You select to launch the Connection Management application. SAP Data Intelligence Cloud opens the Connection Management Application.

3.2.2 Start or Restart a Stable Application Instance

Start or Restart any stable instance of an application in the System Management *Applications* page.

Prerequisites

You must have permission to restart the application.

Note

If the application is nested, start or restart the core (parent) application.

Context


When you open the *Applications* page, all applicable applications appear in rows on the page. At the end of each row, the *Action* column contains action icons.


Procedure

1. Locate the row that contains the application in the *Applications* page.

→ Tip

Use the search feature if you can't find the application by scanning the list.

2. To start the application, choose  (*Start*).
A message appears indicating that the application is started. The value in the Status column changes to "ready".

3. To restart the application, choose  (*Restart*).
A message appears that indicates that the application is restarting. Check the value in the Status column. At first, the status shows “restarting”. The application is restarted when “ready” appears in the Status column. A status of “not started” indicates the application can be restarted.

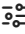

3.2.3 Set Application Configurations

Tenant Administrators can access and change settings in the [Application Configuration and Secrets](#) page.

Context

For descriptions of each general application configuration parameter, see [System Management Application Configuration and Secrets Properties \[page 81\]](#).

Procedure

1. Choose  (*View Application Configuration and Secrets*) in the *Applications* page.
The *Application Configuration and Secrets* page opens.
2. Open the *General* page.
3. Choose  (*Edit*).
4. Change the applicable parameter.
5. Choose *Save*.

3.2.4 Add an Application Secret

Add an application secret to applications for another level of security.

Prerequisites

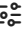

To perform this task, you must be a tenant administrator. You must also have an existing secret file.

Context

The following lists some information about secrets in SAP Data Intelligence Cloud:

- You can't share secrets across tenants.
- SAP Data Intelligence Cloud protects some application secrets and disables the secrets for modification.
- Protected secrets can be modified only by the cluster administrator.

Procedure

1. Choose  ([View Application Configuration and Secrets](#)) in the [Applications](#) page.
2. Choose  ([Add "Secret"](#)).
3. Enter a name for the secret.
4. Browse for and select the applicable secret file.
5. Choose [Create](#).

3.2.5 SAP Data Intelligence Cloud Nested Applications

SAP Data Intelligence Cloud nests applications that consume fewer resources and groups them on a single pod. For example, SAP Data Intelligence Cloud Connection Management and Metadata Explorer are nested applications.

SAP Data Intelligence Cloud manages the routing of nested applications internally by a core application, and always maps nested applications on a stable instance.

When you launch any nested application and there's no stable instance of the application that exists, SAP Data Intelligence Cloud creates a stable instance in the core application. This application consumes resources from a single pod where the core application resides. You can't create new instances of nested applications.

When you delete a stable instance of the core application, SAP Data Intelligence Cloud removes all the instances of the nested applications managed by the core application.

3.2.6 Manage Metadata Explorer Catalog and Data Preparation Memory

Tenant administrators can change the default memory size reserved for the Metadata Explorer catalog and data preparation memory usage.

Context

The following table describes the memory setting for the Metadata Explorer catalog and data preparation application, and it lists the default memory usage limits.

Application	Description	Memory limits
Metadata Explorer catalog	Specifies the amount of memory that the catalog can use to store data in the SAP HANA database memory.	8192 MB
Preparation memory	Specifies the amount of SAP HANA database memory allocated to data preparation processes.	4096 MB

To calculate your memory usage requirements, SAP recommends that you divide approximately 60 percent of the available SAP HANA memory usage between the Metadata Explorer catalog and data preparation. The system uses the other 40 percent for other applications, queries, and logging.

→ Remember

Even though the SAP HANA instance is shared between all tenants, the size limit is for a single tenant. Therefore, consider all tenants when you calculate memory usage division.

🔗 Example

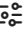

If you have one tenant and a 20-GB database, the size limit is 60 percent divided between the Metadata Explorer catalog and data preparation. But, if you have two tenants, then the space is dropped in half as each tenant could potentially use 60 percent. Therefore, using the calculated memory limit, SAP HANA would run out of memory.

→ Tip

If you run out of memory for the Metadata Explorer catalog, increase the memory usage for the Metadata Explorer catalog and decrease the memory for data preparation.

To adjust the memory usage limits, perform the following steps:

Procedure


1. Choose  ([View Application Configuration and Secrets](#)) in the [Applications](#) page.
2. Open the [General](#) page.
3. Choose  ([Edit](#)).
4. Search for Metadata Explorer catalog and preparation memory individually and edit the limits for each based on your calculation results.

📘 Note

If you set the memory usage limit to 0, then the system sets the value of 8192 MB for Metadata Explorer catalog and 4096 MB for Data Preparation.

For descriptions of the Metadata Explorer catalog and data preparation System Management properties, see [System Management Application Configuration and Secrets Properties \[page 81\]](#).

5. Choose [Update](#).

6. If the Data Application is running, you must recreate it for each user by performing the following substeps:
 - a. In the *Applications* page, search for Data Application.
 - b. Select  (*Restart*) under the **Action** column.

Results

When you open the Metadata Explorer application, the *dataintelligence-app-data* application has the updated memory usage limits shown in the *Memory Usage* tile. For more information about the Metadata Explorer application, see [Using the Metadata Explorer](#).

3.2.7 Manage Metadata Automatic Lineage Extractions

Automatic lineage extraction provides a history of lineage extractions over time and provides information about how SAP Data Intelligence Cloud adds or removes data sources. Automatic lineage extraction also shows which data sources and transformations are included in the output of a data target.

Prerequisites

You must be a tenant administrator to perform this task.

Context

Run automatic lineage extraction on Modeler graphs and Metadata Explorer data preparations.

Caution

There is a limit of 500 graphs in the garbage collector. You should regularly purge old dead graphs that you no longer need to free up space for new graphs to persist. After the limit is reached, any new graphs are immediately garbage collected upon completion and their lineage extraction fails with an error.

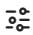

Automatic lineage extraction is turned off by default.

Note

To view the *Automatic Lineage* tab in Metadata Explorer, you must have the `app.datahub-app-data.administrator` policy assignment in System Management. For more information about policies, see [Manage Policies \[page 50\]](#).


To manage automatic lineage properties, perform the following steps:

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad.
2. Open the [Applications](#) page and choose  ([View Application Configuration and Secrets](#)).
3. Choose  ([Edit](#)).
4. Search for Metadata Explorer.

Set the properties for Metadata Explorer lineage as described in the following table.

Metadata Explorer Property	Description
Automatic lineage extraction of Modeler graphs	Select one of the following options: <ul style="list-style-type: none">• enabled_and_publish_datasets: extracts lineage and publishes the dataset to the catalog in the Metadata Explorer. Access the lineage by browsing the connection or the catalog.• enabled_and_do_not_publish_datasets: extracts lineage but doesn't publish the dataset to the catalog. Access the lineage by browsing the connection in the Metadata Explorer.• disabled: turns off the property.
Automatic lineage extraction of Data Preparations	
Days until deletion of automatic lineage extraction from the monitoring task list and catalog	Determines the number of days that Metadata Explorer keeps the automatic lineage extractions on the monitoring list and stored in the catalog. Set to -1 to keep all automatic lineage extractions.
Automatic lineage extraction frequency (in min)	Determines the frequency in which Metadata Explorer monitors the automatic lineage extraction tasks and updates graph information.

5. Choose [Update](#) and choose [OK](#) to close the restart message.
6. Choose [Applications](#).
7. Find the row for [Data App Daemon](#), and choose  ([Restart](#)) at the end of the row.

3.2.8 Manage Size Limit for File Upload and Download

You can change the default size limit of 100 MB for files that are uploaded and downloaded in Metadata Explorer.

Prerequisites

You must be a tenant administrator to perform this task.

Note

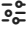

This procedure doesn't change the limits in Data Preparation.

Context


This setting affects all users in the tenant.

To increase or decrease the size limit for file upload and download, perform the following steps:

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad.
2. Open the [Applications](#) page and choose  ([View Application Configuration and Secrets](#)).
3. Open the [General](#) page and choose  ([Edit](#)).
4. Search for [Applications: Size limit on the files that can be uploaded or downloaded \(MB\)](#).
5. Change the size in MB from the default of 100.

If you enter 0, the system defaults to the value of 100 MB.

6. Choose [Update](#) and choose [OK](#) to close the restart message.
7. Open the [Applications](#) page and search for [Data Application](#).
8. Choose  ([Restart](#)) at the end of the row.




3.2.9 Manage Metadata Rule Validation CSV Length

When the set maximum string length of 5000 characters is insufficient, SAP Data Intelligence Cloud issues errors while validating your rules in CSV files.

Context

To manage the string length for rulebooks, perform the following steps:

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad.
2. Open the *Applications* page and choose  (*View Application Configuration and Secrets*).
3. Open the *General* page and choose  (*Edit*).
4. Search for *Maximum CSV column length for rule validation*
5. Adjust the string length as applicable.
6. Choose *Update* and choose *OK* to close the restart message.
7. Open the *Applications* page and search for *Data Application*.
8. Choose  (*Restart*) at the end of the row.

3.2.10 Manage Metadata Failed Records

When records fail a rule, you can save the records to a separate location and examine the failed records later to determine why they failed the rule.

Prerequisites

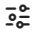

You must have the permission to create, insert, and alter the failed record schema to perform this task. Also gather the following information:

- SAP HANA_DB connection ID.
- Failed data schema name.


Context

To configure SAP Data Intelligence Cloud to save failed records, perform the following steps:

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad.
2. Open the *Applications* page and choose  (*View Application Configuration and Secrets*).
3. Open the *General* page and choose  (*Edit*).
4. Search for “failed record”.
5. Set the following properties based on the descriptions in the table.

Property	Setting
Metadata Explorer: Failed record connection ID, only HANA_DB connection types supported	Enter the SAP HANA_DB connection ID.
Metadata Explorer: Failed record schema, example: /Failed_Records	Enter the name of the schema where the failed data is loaded.

6. Choose *Update* and choose *OK* to close the restart message.
7. Open the Applications page and search for Data Application.
8. Choose  (*Restart*) at the end of the Data Application row.

Related Information

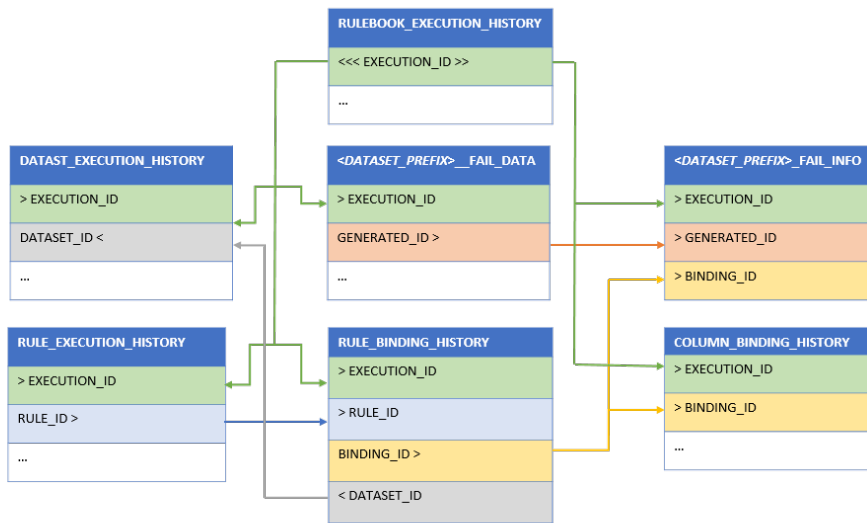
- [Extract Details from Failed Records \[page 38\]](#)
- [Details of the RULEBOOK_EXECUTION_HISTORY Table \[page 39\]](#)
- [Details for the RULE_EXECUTION_HISTORY Table \[page 40\]](#)
- [Details for the DATASET_EXECUTION_HISTORY Table \[page 40\]](#)
- [Details of the RULE_BINDING_HISTORY Table \[page 41\]](#)
- [Details for the COLUMN_BINDING_HISTORY Table \[page 41\]](#)
- [Details of the DATASET_PREFIX_FAIL_DATA Table \[page 42\]](#)
- [Details of the DATASET_PREFIX_FAIL_INFO table \[page 42\]](#)

3.2.10.1 Extract Details from Failed Records

Use your database management system to query metadata tables that contain failed data information.

In Metadata Explorer, save all failed records to a separate table while running a rulebook. During processing, the Metadata Explorer adds metadata tables to the relational database system. Extract information from those tables in your database management system by using SQL-like queries. For example, use a query to find all the rules a given record failed in one run.

The following diagram shows the tables related to the failed data output and how some columns are connected to each table as primary keys.



3.2.10.2 Details of the RULEBOOK_EXECUTION_HISTORY Table

The RULEBOOK_EXECUTION_HISTORY table contains a new record (row) for every rulebook that is run.

The following table describes each column in the RULEBOOK_EXECUTION_HISTORY table.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for the run. This column is the primary key for this table.
RULEBOOK_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for a rulebook.
RULEBOOK_NAME	STRING(256)/ NVARCHAR(256)	User-entered name of the rulebook.
RULEBOOK_DESCRIPTION	STRING(5000)/ NVARCHAR(5000)	User-entered description of the rulebook.
START_TIME	TIMESTAMP	Starting time of the rulebook run.
END_TIME	TIMESTAMP	Ending time of the rulebook run.
STATUS	STRING(20)/ NVARCHAR(20)	Status of the rulebook execution, for example, OK or ERROR.

3.2.10.3 Details for the RULE_EXECUTION_HISTORY Table

The RULE_EXECUTION_HISTORY table contains a new record (row) for every rule that is part of a rulebook run.

The following table contains a description of each column in the RULE_EXECUTION_HISTORY table.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for the run. This column is part of the primary key for this table.
RULE_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for a rule. This column is part of the primary key for this table.
RULE_ID_NAME	STRING(256)/ NVARCHAR(256)	User-entered ID of the rule.
RULE_DISPLAY_NAME	STRING(256)/ NVARCHAR(256)	User-entered name of the rule.
RULE_DESCRIPTION	STRING(5000)/ NVARCHAR(5000)	User-entered description of the rule.
CATEGORY	STRING(256)/ NVARCHAR(256)	Category where the rule is included.

3.2.10.4 Details for the DATASET_EXECUTION_HISTORY Table

The DATASET_EXECUTION_HISTORY table contains a new record (row) for every unique data set that is part of a rulebook run.

The following table describes each column in the DATASET_EXECUTION_HISTORY table.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for the run. This column is part of the primary key for this table.
DATASET_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for a dataset. This column is part of the primary key for this table.
CONNECTION_ID	STRING(256)/ NVARCHAR(256)	Unique connection identifier for the connection used. This ID is what is shown in Connection Management.
QUALIFIED_NAME	STRING(256)/ NVARCHAR(256)	Qualified name of the dataset.
DATASET_PREFIX	STRING(32)/ NVARCHAR(32)	Value for <DATASET_PREFIX> for this run that specifies the <DATASET_PREFIX>_FAIL_DATA and <DATASET_PREFIX>_FAIL_INFO tables.
START_TIME	TIMESTAMP	Starting time of the dataset run.
END_TIME	TIMESTAMP	Ending time of the dataset run.

Column Name	Data Type	Description
STATUS	STRING(20)/ NVARCHAR(20)	Status of the rulebook run, for example, OK or ERROR.
TOTAL_ROWS	FLOATING(8)/DOUBLE	Number of rows that were part of the dataset run.

3.2.10.5 Details of the RULE_BINDING_HISTORY Table

The RULE_BINDING_HISTORY table contains a new record (row) for every unique binding of a rule that is part of a rulebook run.

The following table describes the columns in the RULE_BINDING_HISTORY table.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for the run. This column is part of the primary key for this table.
RULE_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for a rule. This column is part of the primary key for this table.
BINDING_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for a rule binding. This column is part of the primary key for this table.
DATASET_ID	STRING(32)/ NVARCHAR(32)	Unique identifier for a dataset. This column is part of the primary key for this table.

3.2.10.6 Details for the COLUMN_BINDING_HISTORY Table

The COLUMN_BINDING_HISTORY table contains a new record (row) for every unique column mapping that is part of a rule binding. If a rule has multiple parameters, then there are multiple rows in this table, one for each parameter within the rule.

The following table describes the columns in the COLUMN_BINDING_HISTORY table.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/NVARCHAR(32)	Unique identifier for the run. This column is part of the primary key for this table.
BINDING_ID	STRING(32)/NVARCHAR(32)	Unique identifier for a rule binding. This column is part of the primary key for this table.
PARAMETER	STRING(256)/ NVARCHAR(256)	Rule parameter name. This column is part of the primary key for this table.
COLUMN_NAME	STRING(256)/ NVARCHAR(256)	Data source column name to which the specified parameter is bound.

3.2.10.7 Details of the <DATASET_PREFIX>_FAIL_DATA Table

The <DATASET_PREFIX>_FAIL_DATA table contains a new record for every record that fails a run.

The following table describes the columns in the <DATASET_PREFIX>_FAIL_DATA table.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/NVARCHAR(32)	Unique execution identifier that is the same for every record from a single run. This column is part of the primary key for this table.
GENERATED_ID	FLOATING(8)/DOUBLE	Generated row identifier that is created as part of rule processing that uniquely identifies each record. This ID is helpful for linking with the detailed failure information in the <DATASET_PREFIX>_FAIL_INFO table. This column is part of the primary key for this table.
<SOURCE COLUMNS...>	N/A	Contains all columns bound to the rules and can contain additional columns. If a rule expression contains either the <code>is_unique</code> or <code>is_data_dependent</code> function, Metadata Explorer creates an additional column for each function. If this record failed the <code>is_unique</code> or <code>is_data_dependent</code> functions, the additional column contains the Group ID. The additional column name is <code>GRP_ID_<bindingId>_n</code> where <code>n</code> begins at 1 and increments for each function in the rule.

3.2.10.8 Details of the <DATASET_PREFIX>_FAIL_INFO table

The <DATASET_PREFIX>_FAIL_INFO table contains a record for every rule failure.

Because a single record being processed can fail multiple rules, the <DATASET_PREFIX>_FAIL_INFO table can contain multiple records for each failed record. Because a rule may also contain multiple parameters and columns used, the <DATASET_PREFIX>_FAIL_INFO table may also contain multiple records for that case as well.

Column Name	Data Type	Description
EXECUTION_ID	STRING(32)/NVARCHAR(32)	Unique identifier for the run. This column is part of the primary key for this table.

Column Name	Data Type	Description
GENERATED_ID	FLOATING(8)/DOUBLE	<p>Generated row identifier that Metadata Explorer creates as part of rule processing that uniquely identifies each record.</p> <p>This ID is helpful for linking with the detailed failure information in the <code><DATASET_PREFIX>_FAIL_DATA</code> table.</p> <p>This column is part of the primary key for this table.</p>
BINDING_ID	STRING(32)/NVARCHAR(32)	<p>Unique identifier for a rule binding that Metadata Explorer already provides as part of the rule configuration.</p> <p>This column is part of the primary key for this table.</p>

3.2.11 Manage Modeler Customer Registry Secret

Create a secret file containing container registry credentials using information from Docker registries. The Docker registry contains a supported authentication method, such as from AWS, GCP, or Azure. Then create a Dockerfile for graphs where SAP Data Intelligence pulls the images from the external Docker registry.

Prerequisites

You must be a tenant administrator to perform these steps. Also, familiarize yourself with how to create an application secret in System Management by reading [Add an Application Secret \[page 31\]](#).

Procedure

1. Prepare a secret file and name it `vflow-customer-registry.yaml` based on the steps in the following table for your database.

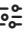

The secret file contains the container registry credentials and the registry host name.

Note

SAP doesn't support authentication schemes or configurations other than the schemes listed in the table.

Database Authentication Method	Steps
AWS	<p>Complete the following steps to create a secret file that contains the authentication method for AWS (Amazon Web Services):</p> <ol style="list-style-type: none"> 1. Fetch the credentials according to the instruction in the Amazon documentation Private Registry Authentication . 2. Use the following code as the System Management secret: <pre data-bbox="874 600 1402 808"> ```yaml - username: "\$AWS_REGISTRY_ACCOUNT_ID" password: "\$AWS_REGISTRY_SECRET" address: "\$AWS_REGISTRY_HOST" ``` </pre>
GCP	<p>Complete the following steps to create a secret file that contains the authentication method for GCP (Google Cloud Platform):</p> <ol style="list-style-type: none"> 1. Fetch the credentials according to the instructions in the Google documentation Service account key . 2. Transform the credentials into a single-line format and use single-quotes for the password field of the secret. 3. Use the following code as the System Management secret: <pre data-bbox="874 1167 1402 1397"> ```yaml - username: "_json_key" password: `\$_{SINGLE_LINE_KEY_FILE_CONTENT}` address: "eu.gcr.io" (or gcr.io, us.gcr.io, or asia.gcr.io) ``` </pre>
Azure	<p>Complete the following steps to create a secret file that contains the authentication method for Azure:</p> <ol style="list-style-type: none"> 1. Fetch the credentials according to the instructions in the Microsoft documentation Authenticate with the service principle . 2. Use the following code as the System Management secret: <pre data-bbox="874 1659 1402 1821"> ```yaml - username: "\$SP_APP_ID" password: "\$SP_PASSWD" address: "***.azurecr.io" ``` </pre>

Database Authentication Method	Steps
Native Docker registry	Use the following code as the System Management secret: <pre> ` ``yaml - username: "\$LOGIN_USERNAME" password: "\$LOGIN_PASSWD" address: "\$ADDRESS" ` `` </pre>

2. Create a new secret in System Management and perform the following specific steps.
 - a. Enter **vflow-customer-registry** for *Name*.
 - b. Choose *Browse* and choose the secret file `vflow-customer-registry.txt`.
 - c. Choose *Create*.
3. Choose  (*Application Configuration & Secrets*).
4. Enter a vSystem secret by performing the following substeps:
 - a. Open the *General* tab and choose  (*Edit*).
 - b. Search for "Modeler: Name" to find the property *Modeler: Name of the vSystem secret containing the credentials for Docker registry*.
 - c. Enter **vflow-customer-registry** for the value.
 - d. Choose *Update* and choose *OK* to close the restart message.

The system takes 5 minutes to reset.

Note

If you change the secret after you initially set the secret, the change happens immediately.

5. Create a Dockerfile that references the registry.

3.3 Managing Users

Tenant administrators can use SAP Data Intelligence Cloud System Management to assign policies to their member users. Member users can use System Management to change their password.

User Types

A user is either a **tenant administrator** or a **member user**. The following table describes the actions allowed for each user type.

User Type	Allowed Actions
Tenant administrator	<p>Performs the following actions within their tenant:</p> <ul style="list-style-type: none"> • Create, view, and delete other tenant administrators and members of their tenant. • Reset their own and other users' passwords. • View and delete user-created instances. • View graphs and the status of all users in their tenant. • Assign policies. • Pause and stop the running graphs for all users in the tenant (using the Monitoring application only). • Restart paused and dead graphs for all users in the tenant (using the Monitoring application only).
Member	Change their own passwords.

User Name and Password

SAP Data Intelligence Cloud integrates with SAP Business Technology Platform User Account and Authentication (SAP BTP UAA). Administrators perform most user management in the SAP BTP UAA. However, you can use System Management to perform the tasks listed in the User Types table based on your user type.

If you're a member user, you must log on to SAP Data Intelligence Cloud for the first time to become visible in System Management. Upon this first login, SAP Data Intelligence Cloud assigns the **sap.dh.member** policy to you automatically. Then your tenant administrator can add to and update your policy assignments.

User Name

You're identified in SAP Data Intelligence Cloud by a unique user name. A user name must have the following characteristics:

- 4 through 64 alphanumeric characters.
- No punctuation marks or white spaces.
- Case insensitive.

Password

By default, the user password must be from 8 through 255 characters, which must include at least one upper-case letter, one lower-case letter, and one digit.

Related Information

[User Policies \[page 47\]](#)

[Create and Delete users and User Instances \[page 47\]](#)

[Change a Password \[page 48\]](#)

[Assign Policies \[page 49\]](#)

[Remove Policies From a User \[page 50\]](#)

3.3.1 User Policies

Tenant administrators use the System Management application to assign policies to their members.

If you're a user member, SAP Data Intelligence Cloud assigns the **sap.dh.member** policy by default when you first log in to the system. The **sap.dh.member** policy provides you with default access to SAP Data Intelligence Cloud System Management and Monitoring applications. Your tenant administrator assigns additional policies based on your role.

The following table lists some of the policies that SAP Data Intelligence Cloud supports. For descriptions of all pre-delivered policies, see [Pre-Delivered Policies \[page 55\]](#).

Policy	Description
sap.dh.admin	Provides access to all applications deployed in your tenant. This policy is assigned only to tenant administrators.
sap.dh.developer	Provides access only to the Modeler application along with the default member access.
sap.dh.metadata	Provides access only to the Connection Management and Metadata Explorer applications along with default member access. <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px; margin-top: 5px;"><p>⚠ Restriction sap.dh.metadata is deprecated and will be removed in future releases. Instead, use app.datahub-app-data.fullAccess as a close replacement. For more information, see Manage Policies [page 50].</p></div>
sap.dh.member	Provides access only to the System Management, Monitoring, and Connection Management applications. This policy is assigned only to tenant members.

3.3.2 Create and Delete users and User Instances

Tenant administrators can create users for their tenants. Additionally, tenant administrators can also remove a user or remove a user's instance.

Prerequisites

Ensure that you're assigned the `sap.dh.admin` policy.

Procedure

1. To create a user, perform the following substeps:
 - a. Open System Management from the SAP Data Intelligence Cloud Launchpad and open the [Users](#) page.
 - b. Choose **+** ([Create User](#)).
 - c. Enter a unique user name.
 - d. Enter the new user's password twice.
 - e. Select [Require password change upon user login](#).

Selecting this option requires the new user to change their password upon their next login.
 - f. Choose [Create](#).
2. Remove a user by performing the following substeps starting in the [Users](#) page:
 - a. Find the applicable user name from the list.
 - a. Choose **✖** ([Delete User](#)) at the end of the row.
 - b. Choose [OK](#) to confirm the deletion.
3. Delete a user's instance by performing the following substeps starting in the [Users](#) page:
In the following substeps, you delete an application instance of a user.
 - a. Choose the applicable user name.

The [Policies](#) page opens.
 - b. Open the [Instances](#) page.
 - c. Find the applicable instance to delete and choose **✖** ([Delete Instance](#)) at the end of the instance row.
 - d. Choose [OK](#) to confirm the deletion.

3.3.3 Change a Password

Tenant Administrators change their own or their tenant members' password. Member users can change their own password.

Context

A user password must be from 8 through 255 characters and must include at least one upper-case letter, one lower-case letter, and one digit.

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad and open the [Users](#) page.
2. Find the applicable user's row.
3. Choose **>** ([Click here to view <user_name> details](#)), at the end of the user's row.

4. Choose [Change Password](#).
5. Enter the new password twice.
6. **Optional:** Select [Require password change upon user login](#).

Selecting this option requires the user to change their password upon their next login. For tenant administrators who create member users' passwords, this step is a security measure to ensure that a user is the only one who knows their password.

7. Choose [Change Password](#).

3.3.4 Assign Policies

Tenant Administrators use the SAP Data Intelligence Cloud System Management application to assign policies to their tenant members.

Context

Policies grant users access to system resources. Tenant administrators use policy assignments for permission control in multiple product components.

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad and open the [Users](#) page.
2. Choose the applicable user name.

The [Policies](#) page opens for the applicable user.

3. Choose [+](#) ([Assign Policy to a User](#)).

A list of eligible policies opens.

4. Choose the policy to assign and choose [Assign](#).

Note

You can add only one policy at a time.

The [Policies](#) page reopens with the new policy listed.

3.3.5 Remove Policies From a User

Tenant administrators can remove policies from tenant member users.

Procedure

1. Open System Management from the SAP Data Intelligence Cloud Launchpad and open the [Users](#) page.
2. Choose the applicable user name.

The [Policies](#) page opens showing the current policies assigned to the user.

3. Find the policy to remove and choose  ([Remove Policy](#)) at the end of the policy row.

3.4 Manage Policies

SAP Data Intelligence Cloud policies grant users access to system resources. Tenant administrators use policies for authorization management to control access to various features and applications. Policies can use multiple attributes, such as user attributes and resource attributes.

Tenant administrators and member users can use Policy Management for the following purposes:

- Tenant administrators can manage policies for all tenant users.
- Member users can manage their user-defined policies.


The [Policies](#) tab in System Management displays the selected user's list of eligible policies. When you select a policy, System Management opens the [Policy Details](#) page, which provides a detailed view of the policy. Details include the policy types and activities allowed plus additional metadata.

Policy ID

Policies are identified with unique policy IDs. There are three categories of policies, distinguished by the policy ID:

- **sap.dh:** predelivered policies, available in every SAP Data Intelligence Cloud tenant.
- **app.APP_NAME:** application-imported policies, available where the application is installed.
- **User-defined policies:** created by users with a user-assigned policy ID.

User-defined Policies

Member users can modify or delete only their user-defined policies. User-defined policies include an [Edit](#) button in the [Policy Details](#) page, and a  ([Delete](#)) icon in the [Policies](#) page.

The following table describes the options when you define a custom policy.

Option	Description
<i>Policy ID</i>	<p>Required. A unique identifier for the policy. Policy IDs are used in assignments.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>When you create a Policy ID, avoid using reserved names that start with "sap" and "app".</p> </div>
<i>Description</i>	Optional. A meaningful description of the policy. The description is visible in the <i>Policies</i> and <i>Users</i> pages of the System Management application.
<i>Add From Policy</i>	Optional. Choose other policies to include (nest) with the new policy. The new policy contains the sum of permissions from all included policies.
<i>Exposed</i>	Required. Indicates whether the policy can be assigned (exposed) to other users. When exposed, tenant administrators can assign the policy to other users. If the policy isn't exposed, it can be used only inside other policies using <i>Add From Policy</i> .
<i>Resources</i>	Optional. The resource for which the policy grants permission.

Assignable Predelivered Policies

There are policies that are assignable directly to users. The following table lists the assignable policies and the policy category. For descriptions of the predelivered policies, see

Category	Policies
Base Policies: policies to which new users in System Management must be assigned.	<p>Base policies include the following:</p> <ul style="list-style-type: none"> sap.dh.clusterAdmin sap.dh.admin sap.dh.member
Additional Exposed Policies: exposed policies that tenant administrators can assign to their tenant users.	<p>Additional exposed policies include the following:</p> <ul style="list-style-type: none"> sap.dh.metadata sap.dh.developer

The following list includes the predelivered policies that tenant administrators can assign directly to member users:

- app.datahub-app-data.fullAccess

Note

Replacement for sap.dh.metadata policy

- sap.dh.admin
- sap.dh.certificate.manage
- sap.dh.clusterAdmin
- sap.dh.connectionContentAllManage
- sap.dh.connectionContentOwnerManage
- sap.dh.connectionCredentialsUnmasked
- sap.dh.connectionsAllRead
- sap.dh.developer
- sap.dh.member

📘 Note

You need to add to an existing policy or create a new policy to gain full access to the features in Metadata Explorer.

🔗 Example

You create a rules dashboard under the following conditions:

- Your users are assigned the **sap.dh.admin** policy without being assigned the **sap.dh.metadata** policy.
- Your users have custom policies that reference **sap.dh.applicationAllStart** or **sap.dh.metadataStart** without being assigned the sap.dh.metadata policy.

To be able to access the Metadata Explorer features fully, add the **sap.dh.metadata** and **app.datahub-app-data.metadataUser** policies, or create a custom policy that includes **app.datahub-app-data.qualityDashboard.manage**.

📘 Note

sap.dh.metadata and **sap.dh.metadataStart** are deprecated and will be removed in future releases. Instead, you can use **app.datahub-app-data.fullAccess** as a close replacement.

Related Information

[Policy Resources and Resource Types \[page 53\]](#)

[Pre-Delivered Policies \[page 55\]](#)

[Nested Policies \[page 58\]](#)

[Application Start Policies \[page 60\]](#)

[Resource Quotas \[page 64\]](#)

[Create a Policy \[page 65\]](#)

[Assign a Policy to a User \[page 67\]](#)

[Mapping Policies to Identity Providers \[page 67\]](#)

Administering Network Policies

[Pre-Delivered Policies \[page 55\]](#)

3.4.1 Policy Resources and Resource Types

A policy resource describes a permission granted on an SAP Data Intelligence system resource. A resource type defines how resource instances can be identified, what activities can be performed on a resource, and how permissions are checked.

Most policy resources consist of the following information:

- **Resource Type:** the entity that the policy protects, such as a connection instance.
- **Activity:** the actions allowed by the policy.
- **Additional metadata:** additional information based on the resource type.

Resource Types

Resource types define how to identify resource instances, what activities you can perform on a resource, and rules for how to check the authorization. Rules are given in a “rego” file that is evaluated by the open policy agent engine.

You can see the resources you’ve recently defined and the inherited resources that are nested under other policies.

The only resource type available is Connection. When an administrator adds the connection resource type to a policy, the administrator chooses the action: read, write, or both. The policy user has rights for the respective actions depending on the underlying identifier of resources.

The following lists resource types in SAP Data Intelligence:

- Application
- systemManagement
- connection
- connectionContent
- Certificate

The following sections contain details about all resource types that are pre-delivered with SAP Data Intelligence. Additional resource types are in a tenant if the SAP Data Intelligence application adds them.

Application

The application resource type controls access to SAP Data Intelligence applications.

The following table describes the characteristics of the application resource type.

Activity	Description	Additional Restriction
start	Allows the user to start the application.	The name of the application limits the permission to a specific application name. The name can include wildcard characters to allow access to multiple applications (GLOB expressions).

systemManagement

The systemManagement resource type allows access to SAP Data Intelligence APIs. The following table describes the characteristics of the systemManagement resource type.

Activity	Description	Additional Restriction
read	Allows user access.	None
write	Allows administrator access.	None

connection, connectionContent

The connection and connectionContent resource types allow access to connections.

The following table describes the characteristics of the connection resource type.

connection

Activity	Description	Additional Restriction
read	Allows showing the connection details.	The <i>Identifier of the connection</i> limits permission to a specific connection ID.
write	Allows modifying the connection.	
ownerRead	Allows showing the connection details for connections created by the user.	The name can be a GLOB expression to allow access to multiple connections.
ownerWrite	Allows modifying connections created by the user.	

The following table describes the characteristics of the connectionContent resource type.

connectionContent

Activity	Description	Additional Restriction
manage	Allows reading from the remote connection end.	The <i>Identifier of the connection</i> limits permission to a specific connection ID.
ownerManage	Allows writing to the remote connection end.	

Certificate

The certificate resource type allows access to certificates for connection to remote systems. The following table describes the characteristics of the certificate resource type.

Activity	Description	Additional Restriction
manage	Allows viewing and modifying certificates.	None

ResourceQuotas

Resource Quotas

ResourceQuotas set resource limits. The following table describes ResourceQuota configurations and values.

Configuration	Possible Values
Resource	CPU, memory, podcount
Resource Limit	Integer count
Target	applications, workloads

For more information, see [Resource Quotas \[page 64\]](#).

3.4.2 Pre-Delivered Policies

Pre-delivered policies are created by default with the tenant.

Some of the pre-delivered application start policies are accessible only in the system tenant or in the customer tenants, depending on the applications that are installed on those tenants.

The following table describes all the pre-delivered policies:

Policy ID	Description	Exposure Flag
app.datahub-app-data.fullAccess	<p>Access to Metadata Explorer, including resources, activities, and dependencies.</p> <p>Allows you to start the Metadata Explorer and to manage connections.</p> <p>For some roles in the Metadata Explorer, permission to start the Modeler is also required.</p>	true
<div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin: 10px 0;"> <p>Note</p> <p>This policy <code>app.datahub-app-data.fullAccess</code>, is a replacement policy for <code>sap.dh.metadata</code>.</p> </div>		
sap.dh.admin	Administer the tenant. Combines access to user, policy, application management, and connection management for the tenant.	true
sap.dh.applicationAllStart	Start all current applications. Override any other start policies.	false
sap.dh.automl.start	Start AutoML.	false
sap.dh.axinoService.start	Start Axino.	false
sap.dh.certificate.manage	Manage server certificates, such as uploading and deleting.	true

Policy ID	Description	Exposure Flag
sap.dh.connectionContentAllManage	Access content of all connections in the tenant in Metadata Explorer. Modify content on the connection such as delete, upload, rename, and create in the Metadata Explorer.	true
sap.dh.connectionContentOwnerManage	Access content of the connections in Metadata Explorer that belong to the user.	true
sap.dh.connectionCredentialsUnmasked	Full visibility to selected connection fields, such as user names. Explicitly enables tenant administrator or pipeline modeler to see the connections' username for troubleshooting connectivity issues and to know with what credentials users access the data sources.	true
sap.dh.connectionMgtStart	Start Connection Management application.	false
sap.dh.connectionsAllRead	Read access to all connection definitions in the tenant.	true
sap.dh.connectionsAllWrite	Write access to all connection definitions in the tenant.	false
sap.dh.connectionsOwnerRead	Read access to the connection definitions that belong to the user.	false
sap.dh.connectionsOwnerWrite	Write access to the connection definitions that belong to the user.	false
sap.dh.datahubAppAuditlog.start	Start Audit Log Viewer.	false
sap.dh.datahubAppCore.start	Start Core Application.	false
sap.dh.datahubAppDaemon.start	Start Data App Daemon.	false
sap.dh.datahubAppData.start	Start Data Application.	false
sap.dh.datahubAppDatabase.start	Start Data Intelligence Application Database.	false
sap.dh.datahubAppDex.start	Start Customer Data Export.	false
sap.dh.datahubAppLaunchpad.start	Start Launchpad.	false
sap.dh.datahubAppLogging.start	Start Monitoring.	false
sap.dh.datahubAppPolicy.start	Start Policy Management.	false
sap.dh.datahubAppPreparation.start	Start Data Preparation.	false
sap.dh.datahubAppScheduler.start	Start Scheduler.	false
sap.dh.datahubAppSystemManagement.start	Start System Management.	false

Policy ID	Description	Exposure Flag
sap.dh.datahubAppTask.start	Start task Application.	false
sap.dh.dataHubFlowAgent.start	Start Flowagent.	false
sap.dh.developer	Access the Pipeline Modeler. Start the Modeler and manage connections.	true
sap.dh.dspGitServer.start	Start Git Server.	false
sap.dh.jupyter.start	Start Jupyter Lab.	false
sap.dh.licenseManager.start	Start License Management.	false
sap.dh.member	Member of tenant. Access to connection management to configure connections of current user. Requires additional application start policies for running applications.	true
sap.dh.metadata	Start Metadata Explorer and manage connections. For some roles in the Metadata Explorer, permission to start the Pipeline Modeler is also required.	true
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <p><code>sap.dh.metadata</code> is deprecated and will be removed in future releases. Instead, you can use <code>app.datahub-app-data.fullAccess</code> as a close replacement. For more information, see Manage Policies [page 50].</p> </div>		
sap.dh.metadataStart	Start Metadata Explorer.	false
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <p><code>sap.dh.metadata</code> is deprecated and will be removed in future releases. Instead, you can use <code>app.datahub-app-data.fullAccess</code> as a close replacement. For more information, see Manage Policies [page 50].</p> </div>		
sap.dh.metricsExplorer.start	Start Metrics Explorer.	false
sap.dh.mlApi.start	Start ML API.	false
sap.dh.mlApiImporter.start	Start ML API Importer.	false
sap.dh.mlDeploymentApi.start	Start ML Deployment API.	false
sap.dh.mlDmApi.start	Start ML DM API.	false
sap.dh.mlDmApp.start	Start ML Data Manager.	false

Policy ID	Description	Exposure Flag
sap.dh.mlScenarioManager.start	Start ML Scenario Manager.	false
sap.dh.mlTracking.start	Start ML Tracking API.	false
sap.dh.modelerStart	Start Pipeline Modeler.	false
sap.dh.modelerUI.start	Start the Pipeline Modeler UI.	false
sap.dh.monitoring	<p>If you're a developer member user (that is, has the sap.dh.developer and sap.dh.member policies assigned to you) and are also assigned the monitoring policy, you can view other user graph instances in the Monitoring application.</p> <p>Tenant administrators assigned with the monitoring policy can view and edit schedules for graphs from all tenants.</p>	true
sap.dh.resourceplanService.start	Start Resource Plan Service.	false
sap.dh.shared.start	Start Shared.	false
sap.dh.stopAppsForOtherUsers	Stop application instances started by any other user. Stop applications owned by an administrator.	false
sap.dh.systemAccess	Access to System Management API-based policy.	false
sap.dh.systemMgtWrite	Change operations in system management and policy management.	false
sap.dh.trainingService.start	Start Training Service.	false

3.4.3 Nested Policies

A nested policy uses the definitions and resources of other policies.

The following table shows policies and their respective nested policies.

Policy ID	Referred Policies
sap.dh.admin	<ul style="list-style-type: none"> • sap.dh.applicationAllStart • sap.dh.certificate.manage • sap.dh.connectionsAllRead • sap.dh.connectionsAllWrite • sap.dh.stopAppsForOtherUsers • sap.dh.systemAccess • sap.dh.systemMgtWrite

Policy ID	Referred Policies
sap.dh.clusterAdmin	<ul style="list-style-type: none"> sap.dh.applicationAllStart sap.dh.certificate.manage sap.dh.connectionsAllRead sap.dh.connectionsAllWrite sap.dh.stopAppsForOtherUsers sap.dh.systemAccess sap.dh.systemMgtWrite
sap.dh.developer	<ul style="list-style-type: none"> sap.dh.connectionsAllRead sap.dh.connectionMgtStart sap.dh.datahubAppLogging.start sap.dh.jupyter.start sap.dh.modelerStart sap.dh.modelerUI.start
sap.dh.member	<ul style="list-style-type: none"> sap.dh.automl.start sap.dh.axinoService.start sap.dh.connectionContentOwnerManage sap.dh.connectionMgtStart sap.dh.connectionsOwnerRead sap.dh.connectionsOwnerWrite sap.dh.datahubAppCore.start sap.dh.datahubAppDaemon.start sap.dh.datahubAppData.start sap.dh.datahubAppDatabase.start sap.dh.datahubAppLaunchpad.start sap.dh.datahubAppScheduler.start sap.dh.datahubAppSystemManagement.start sap.dh.datahubAppTask.start sap.dh.dataHubFlowAgent.start sap.dh.dspGitServer.start sap.dh.metadataStart sap.dh.metricsExplorer.start sap.dh.mlApi.start sap.dh.mlApiImporter.start sap.dh.mlDeploymentApi.start sap.dh.mlDmApi.start sap.dh.mlDmApp.start sap.dh.mlScenarioManager.start sap.dh.mlTracking.start sap.dh.resourceplanService.start sap.dh.shared.start sap.dh.systemAccess sap.dh.trainingService.start

Policy ID	Referred Policies
sap.dh.metadata <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p>Note</p> <p>sap.dh.metadata and sap.dh.metadataStart are deprecated and will be removed in future releases. Instead, you can use app.datahub-app-data.fullAccess as a close replacement. For more information, see Manage Policies [page 50].</p> </div>	<ul style="list-style-type: none"> • sap.dh.connectionContentAllManage • sap.dh.connectionMgtStart • sap.dh.connectionsAllRead • sap.dh.metadataStart

3.4.4 Application Start Policies

Application start policies allow tenant administrators to specify what applications can be started by individual users. Therefore, application start policies allow better control over resource consumption and ease of use for users, who see only a subset of available applications in their launchpad.

A policy to start an application uses the `application` resource type, with the `start` activity and the `name` field. The `name` field is matched against the application identifier (ID).

Note

The application ID is the ID column that appears while listing the available templates with the command-line client. Application ID is also equivalent to the link to the application, for example, `datahub-app-launchpad`.

Create application start policies using the SAP Data Intelligence System Management Command-Line Client or Policy Management.

All pre-delivered applications have a start policy available to use. The specific policies are not exposed; they can only be nested under other policies and not directly assigned to users.

The `sap.dh.member` policy includes most of the available start policies. To restrict the set of startable applications, tenant administrators must create a new policy and reference the appropriate start policies.

The `sap.dh.applicationAllStart` policy allows all applications to be started, without the need to assign specific start policies. Tenant and cluster administrators have this policy assigned to them.

Start policies are checked when an application is started (that is, when a pod is created). The start policies are not checked when the application is accessed; for example, when a user uses the stable link after the pod has been created. In other words, after a tenant application has been started, any user in possession of the stable link can use it.

Some applications may depend on features of other applications. It is the tenant administrator's responsibility to model the policies so that users have the policy to start the application, as well as the policies to start its dependencies.

Users can stop their user applications, but not the user applications of other users, nor tenant application instances. Tenant administrators can stop tenant applications instances, as well as any user application instance in their tenant. It is directed by the `sap.dh.stopAppsForOtherUsers` policy, which is given to

administrators but not to regular members. Users with the `sap.dh.stopAppsForOtherUsers` policy can stop any instance of an application in their tenant, provided that they can start that application (that is, they have a start policy for that application that applies to them).

The list of available applications always depends on the policies of the user. The cluster administrator can see how many applications their users can start by using the `-t` and `-u` flags of the System Management Command-Line Client, which allow some commands to be executed as another tuple tenant/user.

Note

Start policies do not apply in the system tenant.

Related Information

[Required Application Start Policies \[page 61\]](#)

3.4.4.1 Required Application Start Policies

To start an SAP Data Intelligence application, a user must have access to all of the start policies required for the application.

If your user has the application start policy `sap.dh.member`, you already have access to many of the available application start policies. However, if your system administrator uses custom policies, you must have the following application start policies to use applications:

Application	Required Application Start Policy	Dependent Application	Dependent Application ID
Axino	<code>sap.dh.axinoService.start</code>	Axino	<code>axino-service</code>
	<code>sap.dh.connectionMgtStart</code>	Connection Management	<code>datahub-app-connection</code>
Flowagent	<code>sap.dh.dataHubFlow-Agent.start</code>	Flowagent	<code>data-hub-flow-agent</code>
	<code>sap.dh.axinoService.start</code>	Axino	<code>axino-service</code>
	<code>sap.dh.connectionMgtStart</code>	Connection Management	<code>datahub-app-connection</code>
JupyterLab	<code>sap.dh.jupyter.start</code>	Jupyter Lab	<code>jupyter</code>
	<code>sap.dh.automl.start</code>	AutoML	<code>automl</code>
	<code>sap.dh.connectionMgtStart</code>	Connection Management	<code>datahub-app-connection</code>
	<code>sap.dh.datahubApp-Core.start</code>	Core Application	
	<code>sap.dh.datahubAppDaemon.start</code>	Data App Daemon	

Application	Required Application Start Policy	Dependent Application	Dependent Application ID
	sap.dh.datahubApp-Data.start	Data Application	
	sap.dh.datahubAppData-base.start	Data Intelligence App DB	
	sap.dh.datahubAppLogging.start	Monitoring	datahub-app-logging
	sap.dh.datahubAppLaunchpad.start	Launchpad	datahub-app-launchpad
	sap.dh.dataHubFlow-Agent.start	Flowagent	data-hub-flow-agent
	sap.dh.metadataStart	Metadata Explorer	datahub-app-metadata
	<div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p>Note</p> <p><code>sap.dh.metadataStart</code> is deprecated and will be removed in future releases. Instead, you can use <code>app.datahub-app-data.fullAccess</code> as a close replacement. For more information, see Manage Policies [page 50].</p> </div>		
	sap.dh.mlApi.start	ML API	ml-api
	sap.dh.mlDmApi.start	ML DM API	ml-dm-api
	sap.dh.mlScenarioManager.start	ML Scenario Manager	ml-scenario-manager
	sap.dh.mlTracking.start	ML Tracking API	ml-tracking
	sap.dh.modelerStart	Pipeline Modeler	pipeline-modeler
	sap.dh.shared.start	Shared	shared
	sap.dh.trainingService.start	Training Service	training-service
Launchpad	sap.dh.datahubAppLaunchpad.start	Launchpad	datahub-app-launchpad
	sap.dh.shared.start	Shared	shared
Metadata Explorer	sap.dh.metadataStart	Metadata Explorer	datahub-app-metadata
	sap.dh.connectionMgtStart	Connection Management	datahub-app-connection

Application	Required Application Start Policy	Dependent Application	Dependent Application ID
	sap.dh.modelerStart	Pipeline Modeler	pipeline-modeler
	<div style="border: 1px solid #ccc; padding: 5px; background-color: #f9f9f9;"> <p>Note</p> <p>sap.dh.metadataStart is deprecated and will be removed in future releases. Instead, you can use app.datahub-app-data.fullAccess as a close replacement. For more information, see Manage Policies [page 50].</p> </div>		
	sap.dh.datahubAppScheduler.start	Scheduler	datahub-app-scheduler
	sap.dh.modelerUI.start	Modeler	modeler-ui
Metrics Explorer	sap.dh.metricsExplorer.start	Metrics Explorer	metrics-explorer
	sap.dh.automl.start	AutoML	automl
	sap.dh.connectionMgtStart	Connection Management	datahub-app-connection
	sap.dh.jupyter.start	Jupyter Lab	jupyter
	sap.dh.mlApi.start	ML API	ml-api
	sap.dh.mlScenarioManager.start	ML Scenario Manager	ml-scenario-manager
	sap.dh.mlTracking.start	ML Tracking API	ml-tracking
	sap.dh.modelerStart	Pipeline Modeler	pipeline-modeler
	sap.dh.shared.start	Shared	shared
ML API	sap.dh.mlApi.start	ML API	ml-api
	sap.dh.connectionMgtStart	Connection Management	datahub-app-connection
	sap.dh.mlTracking.start	Tracking API	ml-tracking
	sap.dh.modelerStart	Pipeline Modeler	pipeline-modeler
	sap.dh.resourceplanService.start	Resource Plan API	resourceplan-service
	sap.dh.trainingService.start	Training Service	training-service
ML Deployment API	sap.dh.resourceplanService.start	Resource Plan API	resourceplan-service
	sap.dh.trainingService.start	Training Service	training-service
ML Scenario Manager	sap.dh.mlScenarioManager.start	ML Scenario Manager	ml-scenario-manager

Application	Required Application Start Policy	Dependent Application	Dependent Application ID
	sap.dh.connectionMgtStart	Connection Management	datahub-app-connection
	sap.dh.jupyter.start	Jupyter Lab	jupyter
	sap.dh.mlApi.start	ML API	ml-api
	sap.dh.mlDmApi.start	Data Manager	ml-dm-app
	sap.dh.mlTracking.start	Tracking API	ml-tracking
	sap.dh.modelerStart	Pipeline Modeler	pipeline-modeler
	sap.dh.resourceplanService.start	Resource Plan API	resourceplan-service
	sap.dh.shared.start	Shared	shared
	sap.dh.trainingService.start	Training Service	training-service
	sap.dh.modelerUI.start	Modeler	modeler-ui
ML Tracking API	sap.dh.mlTracking.start	ML Tracking API	ml-tracking
Resource Plan Service	sap.dh.resourceplanService.start	Resource Plan Service	resourceplan-service
Shared	sap.dh.shared.start	Shared	shared
System Management	sap.dh.datahubAppSystemManagement.start	System Management	datahub-app-system-management
	sap.dh.rms.start	RMS	rms
	sap.dh.shared.start	Shared	shared
Training Service	sap.dh.trainingService.start	Training Service	training-service
	sap.dh.resourceplanService.start	Resource Plan Service	resourceplan-service

3.4.5 Resource Quotas

You can use resource quotas to restrict the cluster usage for users.

Resource Quotas

Resource quotas are defined as user policies that can be assigned by cluster and tenant administrators. The following is a form of resource quotas policy.

Resource Type	Resource	Resource Limit	Target
resourceQuotas	CPU, Memory, PodCount	Limit (numerical value)	Applications, Workloads

In the resource quotas policy you must specify which resource should be restricted. The resource can be CPU usage, memory usage, or the number of pods that a user is allowed to use. The limit represents the corresponding unit, for CPU in millicpu, for Memory in bytes, and for PodCount, the number of pods. You must also specify the target which should be limited, that is, either applications or workloads.

You can assign multiple resource quotas policies to a user, to restrict a combination of the specified resources. If more than one policy is specified for the same resource, the maximum of the limits in those rules is applied automatically.

If a user wants to schedule an application or a workload that would violate the limits of the assigned resource quotas, the scheduling of the application or workload will be rejected. This happens if the requested resource consumption of the application or the workload in addition to the current resource consumption of the user is higher than the set limit.

Tenant Assignments

Resource quota policies can be assigned to tenants. They then implement restrictions for all users in a tenant. The resource request of a user must be within ones own quota and the request combined with all other user requests must be within the tenant quota.

Example: User A, B, C in tenant T have a CPU quota of 10 millicpu each. The tenant has a CPU quota of 20 millicpu. User A and B launch a 10 millicpu application. User C now cannot launch any application. User C's CPU quota is not touched, but the 20 millicpu tenant quota is already exhausted. When User B stops the application, User C can start one.

Note

Only cluster administrators can set tenant resource quotas.

3.4.6 Create a Policy

Create a policy that grants resource access to users.

Context

Note

The propagation of a change, such as a new assignment, a removal, or a change to the assigned policy, can be delayed until it's visible in the entire cluster. A delay of up to 15 seconds is expected.

Applications that use a permission change directly after a policy assignment has changed, need to retry their requests for this duration.

Procedure

1. Log in to SAP Data Intelligence and choose the *System Management* application from SAP Data Intelligence Launchpad.
2. Select the *Policies* tab.

You can see the list of active policies along with their policy IDs, descriptions, and other information about the policies.

3. To add a policy, click **+**.
4. Fill in the policy details for the policy you want to create.

All required fields are labeled with an asterisk (*).

Note

The policy ID should have a format that contains alphanumeric characters ('a' to 'z'), an underscore (_), and a dot (.). It must not contain any other symbol, and it must not start with 'sap'.

5. Choose whether the policy is exposed or not. Only an exposed policy can be assigned to users.
6. Choose whether the policy is enabled or not. If enabled, the policy resources and those of its referenced policies are available for policy checks. If disabled, assignments of this policy are possible, but have no effect on the user's permissions.
7. (Optional) If you want to copy resources from an existing policy, search for that policy and select it.

The policy and resources are added. Select the *Inherited* tab to see which resources have been added.

8. Click **+** to add your defined resource.

Select a resource type from the available options.

Note


Each policy must have at least one resource added to it, whether from an inherited policy, or a defined resource.

9. Select an allowed activity.

Only one activity is allowed per resource. To define multiple activities for a policy, add additional resources.

10. Define other properties based on the selected resource type.
11. Click *OK*.
12. Click *Create*.

After the policy is created, it is displayed in the policy list.

13. Click  in the *Action* column to *Delete* a policy.
14. To edit a policy, select the policy and click *Edit* in the *Policy Details* page.

3.4.7 Assign a Policy to a User

Use System Management to assign previously created or predefined policies to a user.

Context

ⓘ Note

The propagation of a change, such as a new assignment, a removal, or a change to the assigned policy, can be delayed until it's visible in the entire cluster. A delay of up to 15 seconds is expected.

Applications that use a permission change directly after a policy assignment has changed, need to retry their requests for this duration.

The following steps explain how to add policies to users:

Procedure

1. Log in to SAP Data Intelligence and choose the *System Management* application from the launchpad.
2. Click the *Users* tab on the top of the page.
3. On the left side of the page, select the user or cluster that you'd like to add the policy to.
4. Click the *Policies* tab.
In the *Policies* tab you can see which policies the user already has assigned to them and the description for each policy. You can further view the policy and the defined resources by clicking the **⋮** button and clicking *Manage*.
5. Click **+** and either search for the policy, or select the policy from the list.
6. Click *Assign*.
A message lets you know that the policy was added successfully.
7. To remove a policy, click **⋮** and select *Remove*.

3.4.8 Mapping Policies to Identity Providers

For Identity Provider (IdP) policy users, SAP Data Intelligence administrators can configure automatic policy assignment rules based on data from the identity provider. This is called policy mapping, and maps from a "claim" value in the JWT token of an OpenID Connect IdP to SAP Data Intelligence policies.

Changes in policy assignments from policy mappings always take effect for a new session. If a new policy mapping is configured, users must log in again to see their permission changes. If a policy mapping is changed or deleted and the users are awarded different permissions, the changes do not affect current sessions. The changes only affect sessions after users log in again.

Example of a policy mapping:

ID	Claim	Value	Policies
a3d5	/iss	sap.com	sap.dh.member

Note

The policy mapping feature is only available for OpenID Connect identity providers.

The policy mapping feature can use any JWT Token claims, provided that they are either "simple" data types or lists. For "simple" datatypes (for example, strings or integers), the value of the claim is compared to the policy mapping value as string. For list data types, any list entry that matches the policy mapping value leads to a match. Other data types cannot be compared.

The format of the "claim" is a string following the JSON Pointer specification. The format of the "value" is a string.

By default the policy mapping feature is disabled, it can be configured by changing the policy mapping mode to the **onlyMapping** or **manualAndMapping** (see, Set Policy Mapping Mode section).

Administration Commands

To administer policy mappings, use the SAP Data Intelligence Command-line Client Tool (vctl). You can perform the following:

- List Mappings
- Create a Mapping
- Delete a Mapping
- Set Policy Mapping Mode

Note

The flag `--idp <idp-name>`, for all the policy mapping vctl commands, is optional if there is just one identity provider registered on the tenant. Otherwise is mandatory to specify the identity provider.

List Mappings

To list policy mappings for an IdP, run the following command:

```
vctl idp policy-mapping list --idp <idp-name>
```

Example Output:

ID	Description	Claim	Value
14ade90b-d42f-11eb-a248-06c6b82bae97		/roles	admin
sap.dh.admin, sap.dh.developer			
c7fa7d13-d32d-11eb-a248-06c6b82bae97		/roles	testgroup
sap.dh.member			

In this example, the first mapping assigns two policies (sap.dh.admin and sap.dh.developer) to users that have the claim "roles" with the value "admin". The second mapping assigns sap.dh.member to users that have the

claim "roles" with the value "testgroup". In the setup used here, the "roles" claim corresponds to groups at the identity provider, so the policy mappings assign permissions based on the group membership of the users.

Create a Mapping

To create a mapping, run the following command:

```
vctl idp policy-mapping create --idp <idp-name> "/roles" "admin" "sap.dh.admin"
```

Example Output:

ID	Description	Claim	Value	Policies
0a293f58-d43c-11eb-a248-06c6b82bae97		/roles	admin	
sap.dh.admin				

The `create` command returns the details of the created policy mapping.

Note

Each policy mapping is identifiable by a UUID that is automatically created.

Delete a Mapping

To delete a mapping, run the following command:

```
vctl idp policy-mapping delete <idp-name> 0a293f58-d43c-11eb-a248-06c6b82bae97
```

This command produces no output.

Set Policy Mapping Mode

There are three types of policy mapping modes:

onlyManual: only manual assignment of policies to users by administrators.

onlyMapping: only automatic assignment of policies. Administrators can't assign policies to users.

manualAndMapping: automatic assignment of policies and manual assignment of policies by administrators is allowed.

To set the policy mapping mode, run the following command:

```
vctl idp policy-mapping set-mapping-mode manualAndMapping
```

Example Output:

```
Successfully updated policy mapping mode for idp "xsuaa" to "manualAndMapping"
```

XSUAA User Federation

SAP Data Intelligence integrates with Business Technology Platform (BTP) user management through an "XSUAA" provider, which can be used to login to DI with BTP users. It's possible to register additional identity providers in BTP and have these federated users login to SAP Data Intelligence as well (For more details see, [Configuring External Identity Providers in SAP Data Intelligence \[page 278\]](#)). Using policy mappings, you can

map attributes of these federated users to SAP Data Intelligence policies, as long as you know which claim in the XSUAA access token contains the desired attribute for the federated user. For instance, when using SAP Identity Authentication Service (IAS), you can add users to groups, and when those IAS users login to SAP Data Intelligence you can automatically give them policies by creating a policy mapping that references the claim `/xs.system.attributes/xs.saml.groups`.

The following command will create a policy mapping that assigns the SAP Data Intelligence policy `"sap.dh.admin"` to any IAS users that belong to the `"admin"` group in IAS:

```
vctl idp policy-mapping create /xs.system.attributes/xs.saml.groups admin
sap.dh.admin
```

Related Information

[Manage Policies \[page 50\]](#)

[Configuring External Identity Providers in SAP Data Intelligence \[page 278\]](#)

3.5 Managing Files

SAP Data Intelligence Cloud provides a shared file system for applications that run on the System Management server. The shared file system is synchronized across all application instances in the tenant.

The System Management *Files* page includes the following tabs:

- *My Workspace*: contains files and folders that are accessible only by the current user. All changes that you make in *My Workspace* are visible only to the current user.
- *Union View*: a workspace that groups files and folders from all solutions in the current tenant strategy, along with the files listed in the *My Workspace* tab. The result is a consolidated view of all the files and folders that are accessible to the current user.

By default, System Management provides a tree view of files and folders in both the *My Workspace* and *Union View* tabs.

Note

You can't modify files that come from solutions in the base strategies of the tenant in *Union View*. You can perform read-only operation in *Union View*.

For more information about strategies, files, and permissions, see [Strategies \[page 80\]](#)

Related Information

[Create a Folder \[page 71\]](#)

[Create a File \[page 72\]](#)

- [Delete a File or Folder \[page 72\]](#)
- [Rename a File or Folder \[page 73\]](#)
- [Copy the Path of a File or Folder \[page 73\]](#)
- [Export a File or Folder \[page 74\]](#)
- [Import a File or Solution \[page 74\]](#)
- [Import a Solution from the Solution Repository \[page 75\]](#)
- [Export a File or Folder as a Solution \[page 76\]](#)
- [Export a Solution to Solution Repository \[page 76\]](#)
- [Sharing Files Using Solution Repository \[page 77\]](#)

3.5.1 Create a Folder

Any new folder that you create in the file system are visible only to you. The changes are present only in your workspace.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Choose a file after which the new folder will be located.
3. Choose **+** (*Create new file or folder*) in the toolbar and select *Create Folder*.
4. Enter a name for the new folder.

→ Tip

To create a subfolder, add a forward slash (/) before the name of the new folder.

5. Choose *Create*.

3.5.2 Create a File

Create new files in your workspace.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Choose **+** (*Create new file or folder*) in the toolbar and select *Create File*.
3. Enter a name for the new file.
4. Choose *Create*.

3.5.3 Delete a File or Folder

When you delete a file or folder that is shared with other users, SAP Data Intelligence System Management deletes only the file that you created by creating a white out file for the deleted file in your workspace. To restore the file or folder, delete the white out file in your workspace.

Prerequisites


To perform this task, you must be logged in as a tenant administrator or a member user.

Context

After you delete a file or folder that is shared, other users can continue to access the file in their workspaces.

Procedure

1. Open the *My Workspace* page.
2. Choose the file or folder to delete.

3. Choose  (*Delete selected files or folder*) in the toolbar.
4. Choose *Delete* to confirm the deletion.


3.5.4 Rename a File or Folder

Rename a file or folder in your workspace.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Choose the file or folder to rename and choose  (*Rename file or folder*) in the toolbar.
3. Enter a new name for the file or folder and choose *Rename*.


3.5.5 Copy the Path of a File or Folder

Copy the absolute path of a file or folder to the clipboard.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Select the file or folder to copy.
3. Choose  (*Show Actions*) at the end of the row of the file or folder.
4. Choose *Copy Path*.
The file or folder path is copied to your clipboard.

3.5.6 Export a File or Folder

Export a file or folder to your local system.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Select one or more files or folders.
3. Choose the *Export files or solution* icon in the toolbar and choose *Export files*.
4. **Optional:** Enter a different file or folder name in *File Name*.
5. Choose *Export Files*.

Results

The file appears as a `.tgz` file in your local Downloads folder.

Note

If there's a conflict with the file or solution being imported, choose to either replace the conflicting files or retain the existing files. Currently conflict resolution is supported only in user and union workspace.

3.5.7 Import a File or Solution

You can import a file or a solution as a file. A solution file is a packaged application compressed into a `TAR.GZ` file.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Context

A solution is a packaged application compressed into a Zip file. The Zip file contains a `manifest.json` file with details such as, solution name, version, description, and so on.

Procedure

1. Open the *My Workspace* page.
2. Select the *Import file or solution* icon in the toolbar and select either *Import File* or *Import Solution File*.
3. Browse for and choose the file or solution to upload.
4. Choose *Open*.

Results

Note

When you import a solution, in the respective space, the `/files/manifest.json` file is created or updated with the manifest of the imported solution. In addition, the import functionality doesn't restrict file import based on file extension.

3.5.8 Import a Solution from the Solution Repository

You can import a solution from the solution repository and modify it as applicable.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Select the *Import file or solution* icon in the toolbar and then select *Import Solution from solution repository*.
3. Choose a solution from the existing solutions in the repository.
4. Choose *Import Solution*.

Results

Note

If there's a conflict with the file or solution being imported, you can choose to either replace the conflicting files or retain the existing files. Currently conflict resolution is supported only in user and union workspace.

Note

When you import a solution in the respective space, the `/files/manifest.json` file is created/updated with the manifest of the imported solution.

3.5.9 Export a File or Folder as a Solution

Export a file or folder as a solution to your local system.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the [My Workspace](#) page.
2. Choose one or more files or folders.
3. Choose the [Export files or solution](#) icon in the toolbar and then choose [Export as solution](#).
4. Update the VSolution JSON in the text box with the applicable name, version, format, dependencies, and description as applicable.
5. Choose [Export as Solution](#).

3.5.10 Export a Solution to Solution Repository

Export a solution to the solution repository.

Prerequisites

To perform this task, you must be logged in as a tenant administrator or a member user.

Procedure

1. Open the *My Workspace* page.
2. Choose the applicable solution.
3. Choose the *Export files or solution* icon in the toolbar and then choose *Export as solution to solution repository*.
4. Update the VSolution JSON in the text box and the required name, version, format, dependencies, and description.
5. Choose *Export as Solution*.

3.5.11 Sharing Files Using Solution Repository

Create and modify files in the SAP Data Intelligence System Management application, upload them to the solution repository, and obtain other existing file solutions from the solution repository.

Context

Note

Solutions that contain files created using the operation *Export as solution to solution repository* are visible by all users in the tenant. Solutions in the repository are isolated by tenant; you can share with another tenant by downloading and uploading to another tenant.

The following steps provide an overall workflow:

Procedure

1. Open the System Management application.
2. Open the *Files* tab.
3. Select one or more files, select **⋮** (*Show actions*) and choose *Export as solution to solution repository*.

The *Export as solution to solution repository* dialog box opens showing the manifest.

4. Enter a name for the solution and select *Export as Solution*.
A new solution with the selected files is created in the repository.
5. To import the exported solution using the System Management application, perform the following subtasks:
 - a. Open a new tenant in SAP Data Intelligence and open the System Management application.
 - b. Select **⋮** (*Show actions*) and choose *Import solution from solution repository*.

All solution files are extracted to your current workspace. If conflicts occur while importing files, the Modeler allows you to select files to keep.

6. To share files with other tenants using vctl commands, choose the following commands:

Command	Description
<code>vctl vrep [space] import-solution [name] [version] [destination]</code>	<p>Imports a solution from the repository of [space], to the specified [destination] path.</p> <p>Use the parameters [name] and [version] to import the exact solution.</p> <p>To set the conflict resolution mode, such as handling the file conflicts during the operation, use the flag <code>-r</code>.</p>
<code>vctl vrep [space] export-solution [name] [version] [source...]</code>	<p>Creates a solution with the files of [space] described by [source] and uploads it to the repository.</p> <p>Use the parameters [name] and [version] to export the exact solution.</p> <p>To pass solution dependencies, use the flag <code>-d</code>, if necessary.</p>

3.6 Manage Tenant

As a tenant administrator, you can manage solutions and strategy in the System Management application. You can extend your strategies by importing your own solutions.

Prerequisites

You are a tenant administrator.

Context



SAP Data Intelligence default tenant is assigned with a default strategy, which is inherited from the parent strategy available in the SAP Data Intelligence installation package. You can extend this strategy by importing your own solutions.

Feature	Description
Solutions	A solution is a packaged application compressed into a ZIP file. The ZIP file contains a manifest.json file with details such as, solution name, version, description, etc.

Feature	Description
Strategies	Strategies can reference a parent strategy. This means that all solutions included in the parent strategy are automatically available to the strategies that reference the parent strategy. If the parent strategy is updated, then all the strategies derived from the parent strategy are also updated. For more information, see Strategies [page 80]

Procedure

1. Open the System Management user interface and choose *Tenant* tab.
2. To manage solutions, choose the appropriate option:

Option	Description
Create a solution	<ol style="list-style-type: none"> 1. Choose the <i>Solutions</i> tab. 2. In the <i>Solutions</i> panel, choose the + icon. 3. Browse to select and upload the solution file. 4. Choose <i>Create Solution</i>.
Delete a solution	<ol style="list-style-type: none"> 1. Choose the <i>Solutions</i> tab. 2. In the <i>Solutions</i> panel, choose the  icon.
Download a solution	<ol style="list-style-type: none"> 1. Choose the <i>Solutions</i> tab. 2. In the <i>Solutions</i> panel, choose the  icon.

Note

You cannot download the sap.core solutions.


3. Tenant administrators can update the solutions in the strategy. They can perform only the following actions:
 - Reorder the solutions in the strategy.
 - Add or remove solutions in the strategy.


Restriction

They cannot modify or update the solutions of the parent strategy.

- They can change the inherited parent strategy and reference a different parent strategy.

To manage strategies, choose the appropriate option:

Option	Description
Update a strategy	<ol style="list-style-type: none"> 1. Choose the <i>Strategy</i> tab. 2. Select the strategy. 3. Click  (Edit) icon. 4. Choose the + (Add Solutions) icon. 5. Select the required solutions.

Option	Description
	6. Choose Add .
Update the default strategy.	<p>Upgrading the SAP Data Intelligence instance does not affect the strategies assigned to the tenant. However, after the upgrade is complete, you can consume the new content.</p> <ol style="list-style-type: none"> 1. Choose the Strategy tab. 2. Select the strategy. 3. Click  (Edit) icon. 4. From the Parent Strategy dropdown list, select the new default strategy for your tenant. 5. Choose Save. 6. Activate the changes by restarting the application.

3.7 Strategies

Strategies are cluster-scoped entities used to administer the applications and content available to the tenants.

Strategies work similar to the layering concept of Docker images, producing a filesystem snapshot, where the content of the solutions is layered on top of each other. When a path collision occurs between the content of the included solutions, the resulting view's contents are determined by the ordering of the solutions (the top most prevails).

Strategies are extended by referencing another strategy as a parent strategy, thus inheriting the solutions of its parent. Strategies are divided into the following categories:

- Base strategies: self-contained; no references to other strategies.
- Extension strategies: strategies inheriting from a base strategy.

Depending on the type of strategy assigned to a tenant, the tenant administrator has limited permissions to operations on strategies. The following table describes the operations available to a tenant administrator:

Operation	Base Strategy	Extension Strategy
Create	No	No
Delete	No	No
Get	If base strategy is the base to this tenant's strategy	If extension strategy is assigned to this tenant.
List	If base strategy is the base to this tenant's strategy	If extension strategy is assigned to this tenant.
Modify (Add/Remove/Reorder Solutions)	No	<p>Tenant administrator can modify an extension strategy if:</p> <ul style="list-style-type: none"> • the extension strategy is assigned to their tenant • the extension strategy isn't assigned to other tenants
Set-Parent	No	No

Operation	Base Strategy	Extension Strategy
		Also, tenant administrators can't remove the parent reference.

If you're a tenant administrator and are updating an extension strategy, it's advised that you include the following set of essential solutions.

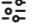

- `vsolution_vsystem_ui-<CURRENT_DATAHUB_VERSION>`: This solution is required for the WebUI.
- `vsolution_shared_ui-<CURRENT_DATAHUB_VERSION>`: This solution is required for the WebUI.
- `certificates-<CURRENT_DATAHUB_VERSION>`: This solution is required to gain access to client certificates created during installation:
- `installer_config-<CURRENT_DATAHUB_VERSION>`: This solution is the default parameter settings for the tenant. For example, Docker registry for Modeler images.

Note

You can't modify files that come from solutions present in the base strategies in Union View. For more information about operations in the Union View workspace, see [Create a Folder \[page 71\]](#).

3.8 System Management Application Configuration and Secrets Properties

Tenant administrators can access and set System Management properties, which control every aspect of SAP Data Intelligence, including areas such as monitoring, connection management, and the Modeler.

To access and edit the System Management application configuration and secrets properties, choose  ([View Application Configuration and Secrets](#)) in the *Applications* tab, and then choose  ([Edit](#)).

Application	Property	Description	Default Setting	Impact	More Information
app-core	hana max pool size defined for CM	Deprecated in SAP Data Intelligence version 2022.09	20	N/A	N/A
app-core	hana pool acquire timeout in millis defined for CM	Deprecated in SAP Data Intelligence version 2022.09	10000	N/A	N/A
app-core	The flag to control features in proxying implementation	Deprecated in SAP Data Intelligence version 2022.09	True (on)	N/A	N/A
app-core	memory resource limit	Protected	1 Gi	N/A	N/A
app-core	memory resource request	Protected	500 Mi	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Applications	Trace Level	Create a trace level log for the category you select from the list.	info	N/A	Options include the following: <ul style="list-style-type: none"> • info • debug • path • warning • error • fatal
Applications	Node V8 Options	Complete with SAP Support's help. Use to create specific Node V8 configurations for the runtime engine.	<code>--max-old-space-size=900</code> <code>--max-http-header-size=16000</code>	N/A	N/A
Applications	Size limit on the files that can be uploaded or downloaded (MB)	Sets the maximum file size for files uploaded to or downloaded from connections in Metadata Explorer.	100 MB	The size limit protects against extended periods of high-network bandwidth. Increase the size limit to support the larger files based on your networks bandwidth.	For more information, see the following topic: Manage Size Limit for File Upload and Download [page 36]

Application	Property	Description	Default Setting	Impact	More Information
Applications	Preparation memory usage limit (MB)	Sets the amount of SAP HANA database memory allocated to data preparation processes.	4096	Controls the amount of memory that the Metadata Explorer Data Preparation app can use in the SAP HANA database.	For more information, see the following topic: Manage Metadata Explorer Catalog and Data Preparation Memory [page 32]

❖ Example

If an SAP HANA instance is shared among multiple tenants or on a low-memory system, set this option accordingly to avoid running out of memory.

Balance this setting with the Metadata catalog memory usage limit (MB).

Application	Property	Description	Default Setting	Impact	More Information
Data Application	Maximum number of concurrent tasks that execute pipelines in Modeler (automatic lineage excluded)	<p>Limits the number of batch processes that run in parallel. Specify a value from -1 to 10.</p> <p>The Modeler checks that the number of concurrent tasks don't exceed the set maximum. It checks for all types of tasks, such as graphs, rulebooks, preparations, profiles, and publications.</p>	-1	<p>The default of -1 removes any limit for the number of concurrent tasks that can run.</p> <p>Multiple processes running concurrently is expensive and slows your system. The higher the setting the greater the impact on system performance.</p>	For more information about how this setting relates to other data application settings, see Configure App-Data [page 264] .
Data Application	The HANA maximum pool timeout in milliseconds (500, 100000) for each nested application	Specifies the time before the system sends a timeout error when the Metadata Explorer establishes a database connection to the SAP HANA application.	5000	<p>Consider increasing the value under the following circumstances:</p> <ul style="list-style-type: none"> You experience application data timeouts. SAP HANA is slow to create a connection. The pool is maxed out. <p>Use this pool timeout setting with the setting for The HANA maximum pool size (1,300) for each nested application.</p>	For more information about how this setting relates to other data application settings, see Configure App-Data [page 264] .

Application	Property	Description	Default Setting	Impact	More Information
Data Application	The HANA maximum pool size (1, 300) for each nested application	Limits the maximum number of connections the application makes to the SAP HANA database at one time.	100	<p>Consider increasing this setting under the following circumstances:</p> <ul style="list-style-type: none"> You have multiple users on Metadata Explorer that are using up the number of connections in a pool. You experience application data timeouts. SAP HANA is slow to create a connection. The pool is maxed out. <p>Use this pool size setting with the setting in <i>The HANA maximum pool timeout in milliseconds (500, 100000) for each nested application</i>.</p>	For more information about how this setting relates to other data application settings, see Configure App-Data [page 264] . Also see the <i>Sizing Guide for SAP Data Intelligence</i> .
Data Application	Time to live after logout (minutes)	Specifies the time before the metadata application pod is stopped to save on resources.	1440	This setting saves resources. If there aren't any active users using Metadata Explorer, the pod is stopped after the set number of minutes.	For more information about configuring application data properties, see Configure App-Data [page 264]

Application	Property	Description	Default Setting	Impact	More Information
Data Application Daemon	The HANA pool timeout in milliseconds (500, 100000) for each nested application	Specifies the time before the system sends a timeout error when the Metadata Explorer Daemon establishes a database connection for the daemon SAP HANA application.	5000	<p>Consider increasing the value under the following circumstances:</p> <ul style="list-style-type: none"> You experience an increase in SAP HANA timeout errors in the daemon logs. You experience application data timeouts. SAP HANA is slow to create a connection. The pool is maxed out. <p>Use this pool timeout setting with the setting for <i>The HANA maximum pool size (1,300) for each nested application</i>.</p>	For more information about how this setting relates to other data application settings, see Configure App-Data [page 264] .

Application	Property	Description	Default Setting	Impact	More Information
Data Application Daemon	The HANA maximum pool size (1,300) for each nested application	Limits the maximum number of connections that the data application daemon makes to the SAP HANA database.	50	<p>Consider increasing the value under the following circumstances:</p> <ul style="list-style-type: none"> You experience an increase in the number of SAP HANA timeout errors in the daemon logs. You experience application data timeouts. SAP HANA is slow to create a connection. The pool is maxed out. <p>Use this pool size setting with the setting for <i>The HANA pool timeout in milliseconds (500, 100000) for each nested application</i>.</p>	For more information about how this setting relates to other data application settings, see Configure App-Data [page 264] .
HANA-client	Use SSL for Hana connection	Because TLS is required for all SAP HANA connection types, don't change the default setting for this property.	True (on)	TLS is required, so don't turn this toggle off.	As of SAP Data Intelligence version 2023.08, TLS is required for all SAP HANA connections.

Application	Property	Description	Default Setting	Impact	More Information
Metadata Explorer	Automatic lineage extraction of Modeler Graphs	<p>Determines whether to extract lineage automatically during Modeler graph (pipeline) execution and optionally publish the lineage information for each dataset to the catalog.</p> <p>Options include the following:</p> <ul style="list-style-type: none"> • enabled_and_publish_datasets • enabled_and_do_not_publish_datasets • disable 	disable	You should regularly purge old dead graphs that you no longer need to free up space for new graphs to persist. After the limit is reached, any new graphs are immediately garbage collected upon completion and their lineage extraction fails with an error.	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Manage Metadata Automatic Lineage Extractions [page 34] • Analyze Data Lineage (Data Governance User Guide) •
Metadata Explorer	Automatic lineage extraction of Data Preparations	<p>Determines whether Metadata Explorer extracts lineage automatically for data preparations and optionally publish the lineage information for each dataset to the catalog.</p> <p>Options include the following:</p> <ul style="list-style-type: none"> • enabled_and_publish_datasets • enabled_and_do_not_publish_datasets • disable 	disable	You should regularly purge old dead graphs that you no longer need to free up space for new graphs to persist. After the limit is reached, any new graphs are immediately garbage collected upon completion and their lineage extraction fails with an error.	<p>For more information, see the following topics:</p> <ul style="list-style-type: none"> • Manage Metadata Automatic Lineage Extractions [page 34] • Analyze Data Lineage (Data Governance User Guide) •

Application	Property	Description	Default Setting	Impact	More Information
Metadata Explorer	Days until deletion of automatic lineage extractions from the monitoring task list and catalog. Set to -1 to keep all extractions.	Determines the number of days that automatic lineage extractions are kept on the monitoring list and stored in the catalog.	-1	The default setting of -1 keeps the lineage extractions on the monitoring list and stored in the catalog indefinitely. To save storage space, set to greater than 0 to automatically delete the catalog records after the set number of days.	For more information, see the following topics: <ul style="list-style-type: none"> • Manage Metadata Automatic Lineage Extractions [page 34] • Analyze Data Lineage (Data Governance User Guide) •
<div style="border-left: 2px solid #0070C0; padding-left: 10px; margin: 10px 0;"> <p>Note</p> <p>You should regularly purge old dead graphs that you no longer need to free up space for new graphs to persist. After the limit is reached, any new graphs are immediately garbage collected upon completion and their lineage extraction fails with an error.</p> </div>					
Metadata Explorer	Automatic lineage extraction frequency (in min)	Determines the frequency in which Metadata Explorer monitors the automatic lineage extraction tasks and updates graph information.	15	You should regularly purge old dead graphs that you no longer need to free up space for new graphs to persist. After the limit is reached, any new graphs are immediately garbage collected upon completion and their lineage extraction fails with an error.	For more information, see the following topics: <ul style="list-style-type: none"> • Manage Metadata Automatic Lineage Extractions [page 34] • Analyze Data Lineage in the <i>Data Governance User Guide</i>

Application	Property	Description	Default Setting	Impact	More Information
Metadata Explorer	Maximum CSV column length for rule validation	Sets the maximum length allowed for a string column. Controls the string length when you process rules on CSV files.	5000	The default of 5000 is large enough for most scenarios. However, increase the value if you have datasets that have longer strings, such as JSON files or corrupted files. For JSON files, consider setting to 8000.	N/A
Metadata Explorer	Failed record connection Id, only HANA_DB connection types supported	Designates the connection ID for the location to store failed record information for rulebook executions. Designate a location to save all failed records, from rulebook executions to an existing SAP HANA database, for later review or processing.	N/A	N/A	N/A
Metadata Explorer	Failed record schema, example: / Failed_Records	Designates the schema in which the failed record tables are created. The location contains the schema of an existing SAP HANA database for storing all failed records for later review or processing.	N/A	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Metadata Explorer	Metadata catalog memory usage limit (MB)	Specifies the amount of memory that the Metadata Explorer catalog can use to store data in the SAP HANA database memory.	8192 MB	Balance this setting with the setting in <i>Preparation memory usage limit (MB)</i> .	For more information, see the following topic: Manage Metadata Explorer Catalog and Data Preparation Memory [page 32]

❁ Example

If an SAP HANA instance is shared among multiple tenants or on a low-memory system, set this option to avoid running out of memory.

Application	Property	Description	Default Setting	Impact	More Information
Metadata Explorer	Resources for Rule Graph Processing	Set a multiplier, such as 125%, 150%, or 200%, to increase the set CPU and memory limits. Allows unusual resource-intensive rulebooks to complete.	100%	<p>The Modeler pre-sets CPU and memory usage limits based on expected resource usage. The default preset values are sufficient for most cases. However, when your graph execution uses more than the preset CPU and memory, the Modeler stops execution and issues an error.</p> <p>The Modeler allows a graph to exceed preset limits when the table being processed is larger than normal.</p> <p>Consider increasing the percentage when you have resource-intensive rulebooks to process, and the set limits for CPU and memory usage aren't enough.</p>	N/A

Application	Property	Description	Default Setting	Impact	More Information
Metadata Explorer	Maximum days to keep scheduled task information; the most recent run is always kept	Specifies the number of days that information from scheduled tasks is kept for review.	365	In Metadata Explorer, a user can schedule publications, rulebooks, and profiling tasks to run on a schedule. For example, a publication can be scheduled to ensure the latest metadata is always accurate and available. After the set number of days, information about a past execution of the scheduled task will be deleted unless that execution is the most recent that has run.	N/A
Axino	Enable auto-scaling for Axino	Auto scaling provides multiple service instances to balance the workload when there's a high number of workload requests.	False (off)	<p>Without sufficient memory, some or all of your graphs can fail because the system runs out of memory. If you plan to have more than 1 Axino pod running at the same time, consider turning on this option.</p> <p>When enabled, also set <i>Maximum number of replicas</i> and <i>Minimum number of replicas</i> to set the number of expected Axino pods.</p> <p>When disabled, use <i>Resources CPU (Limit)</i> and <i>Resources memory (Limit)</i> to set the resource limits for Axino pod.</p>	<p>Axino establishes connections between remote ABAP-based SAP systems and SAP Data Intelligence. Axino enables bi-directional data transfer.</p> <p>For complete information about auto scaling of ABAP Pipeline Engine (Axino), see SAP Note 3148794.</p>

Note

If you turn on this parameter, you must stop all running ABAP-related graphs and then restart the Axino service.

Application	Property	Description	Default Setting	Impact	More Information
Axino	Secret for PSE	<p>Enter the secret for your ABAP Personal Security Environment (PSE).</p> <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;"> <p>⚠ Restriction</p> <p>Use this parameter only when you have issues connecting with ABAP.</p> </div>	N/A	N/A	<p>Axino establishes connections between remote ABAP-based SAP systems and SAP Data Intelligence. Axino enables bi-directional data transfer.</p> <p>For instructions to secure your network communication for ABAP, see Configure SNC for ABAP [page 127].</p> <p>For more information about issues connecting to ABAP, see SAP Note 2849542.</p>
Axino	Enable work process per connection	Protected as of SAP Data Intelligence version 2023.03. Only tenant administrators can see this parameter. Tenant administrators must go through DevOps to change the value.	True (on)	N/A	N/A
Axino	Resources CPU (Limit)	This setting, along with Resources Memory (Limit) controls the resource limit for a single Axino pod. When Enable auto-scaling for Axino isn't on, this setting sets the upper limit for CPU resources.	2 vCPU	<p>If Axino is working efficiently, there's no need to change the default setting of 2 vCPU.</p> <p>SAP doesn't recommend that you enable Enable auto-scaling for Axino and also change the Resource CPU (Limit).</p>	<p>Axino establishes connections between remote ABAP-based SAP systems and SAP Data Intelligence. Axino enables bi-directional data transfer.</p> <p>For more information about this setting, see SAP Note 3148794.</p>

Application	Property	Description	Default Setting	Impact	More Information
Axino	Resources Memory (Limit)	This setting, along with <i>Resources CPU (Limit)</i> controls the resource limit for a single Axino pod. When <i>Enable auto-scaling for Axino</i> isn't on, this setting sets the upper limit for CPU resources.	4 Gi	<p>If Axino is working efficiently, there's no need to change the default setting of 4 Gi.</p> <p>SAP doesn't recommend that you enable <i>Enable auto-scaling for Axino</i> and also change the <i>Resource CPU (Limit)</i>.</p>	<p>Axino establishes connections between remote ABAP-based SAP systems and SAP Data Intelligence. Axino enables bi-directional data transfer.</p> <p>For more information about this setting, see SAP Note 3148794.</p>
Axino	Maximum number of replicas	<p>Sets the maximum number of Axino pods expected to be processed with auto-scaling.</p> <p>Applicable when you enable <i>Enable auto-scaling for Axino</i>.</p>	4	<p>The connection between Axino and ABAP systems is long-polling. If you enable <i>Enable auto-scaling for Axino</i>, SAP suggests that you set the <i>Maximum number of replicas</i> and <i>Minimum number of replicas</i> to the same values, such as 2/2 or 3/3 or 4/4. Setting the minimum and maximum to the same values helps balance the load of Axino pods.</p>	<p>Axino establishes connections between remote ABAP-based SAP systems and SAP Data Intelligence. Axino enables bi-directional data transfer.</p> <p>For more information about this setting, see SAP Note 3148794.</p>

Application	Property	Description	Default Setting	Impact	More Information
Axino	Minimum number of replicas	<p>Sets the minimum number of Axino pods expected to be processed with auto-scaling.</p> <p>Applicable when you enable <i>Enable auto-scaling for Axino</i>.</p>	4	<p>The connection between Axino and ABAP systems is long-polling connections. If you enable <i>Enable auto-scaling for Axino</i>, SAP suggests that you set the <i>Maximum number of replicas</i> and <i>Minimum number of replicas</i> to the same values, such as 2/2 or 3/3 or 4/4. Setting the minimum and maximum to the same values helps balance the load of Axino pods.</p>	<p>Axino establishes connections between remote ABAP-based SAP systems and SAP Data Intelligence. Axino enables bi-directional data transfer.</p> <p>For more information about this setting, see SAP Note 3148794.</p>
Flowagent	Enable ABAP connector which support RFC connections (deprecated)	Deprecated	N/A	N/A	N/A
Flowagent	Connection Management URL (deprecated)	Deprecated	N/A	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Flowagent	File upload chunk size in MB	Sets the maximum chunk size of data to upload using the Flowagent service.	-1	Larger values cause greater memory consumption in the Flowagent application pod. The recommended maximum value is 10 MB for maintaining stability of the Flowagent service. A negative value sets the service to use the internal default value of 5 MB.	N/A

Note

This setting is applicable for following data storage applications:

- S3 (Amazon Simple Storage Service)
- WASB (Windows Azure Storage Blob)
- ADL (Azure Data Lake, deprecated)
- GCS (Google Cloud Storage)

Application	Property	Description	Default Setting	Impact	More Information
				<ul style="list-style-type: none"> HDFS (Hadoop Distributed File System) 	
Flowagent	Number of threads to use for file uploads	<p>Sets the maximum number of threads to use for uploading data using the Flowagent service.</p> <div style="border: 1px solid orange; padding: 5px;"> <p>⚠ Restriction</p> <p>Applicable to WASB (Windows Azure Storage Blob) and S3 (Amazon Simple Storage Service) only.</p> </div>	-1	A higher value causes higher CPU and memory usage in the Flowagent service. Therefore, the maximum value recommended is 4 threads. A negative value sets the service to use the internal default value of 2 threads.	N/A
Flowagent	Instances	Specifies the number of instances for the tenant workload.	1	The maximum setting is 2 instances for the tenant.	N/A

Application	Property	Description	Default Setting	Impact	More Information
Flowagent	Trace Level	<p>Indicates the type of trace level the system reports on. Possible values include:</p> <ul style="list-style-type: none"> • None • Fatal • Error • Warning • Info • Debug 	info	<p>If you change the level, you must restart Flowagent.</p> <div style="border-left: 2px solid orange; padding-left: 10px; margin-top: 10px;"> <p>⚠ Caution</p> <p>Use debug level only for collecting debug information. After collecting debug information, turn off debug level. If you run in debug level, the application can become unstable because of excessive logging.</p> </div>	For descriptions of all trace levels, see Trace Severity Levels .
Git Terminal	Time to live after logout (minutes)	<p>Specifies the time before the system stops the Git Terminal application pod. The time limit saves resources.</p>	1440	<p>If there aren't any active users using the Git Terminal application, the system stops the pod after the set number of minutes.</p>	N/A
Git Terminal	Node V8 Options (bytes)	<p>Specifies the maximum size of HTTP headers.</p> <p>Can be set only by system tenants.</p>	16000	<p>The total size of HTTP headers received by node.js server must not exceed the specified number of bytes.</p>	N/A
Jupyter Lab	Jupyter Lab Directory	<p>The location for your Jupyter lab.</p>	/vhome/dsp	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Jupyter Lab	Jupyter Resources CPU (Limit)	The maximum amount of CPU that can be used by the application.	0.5 CPU unit	N/A	For more information about CPU units, see https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu .
Jupyter Lab	Jupyter Resources Memory (Limit)	The maximum amount of memory that can be used by the application.	4 Gi	N/A	N/A
Jupyter Lab	Jupyter Max Request Body Size	The maximum size for the body of an HTTP request made to the application. This setting also limits the maximum size of files uploaded to Jupyter lab.	10 M	Set to 0 to disable the size validation.	N/A
Jupyter Lab	Jupyter Resources CPU (Request)	The minimum amount of CPU that is reserved for the application. The application can never consume more than the CPU amount indicated.	0.1 CPU unit	N/A	For more information about CPU units, see https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu .
Jupyter Lab	Jupyter Resources Memory (Request)	The minimum amount of memory that is reserved for the application.	1 Gi	N/A	N/A
Launchpad	Launchpad Jupyter Lab Directory	Deprecated in SAP Data Intelligence version 2023.03.	N/A	N/A	N/A
Launchpad	Time to live after logout (minutes)	Specifies the time before the system stops the Launchpad application pod. The time limit saves resources.	1440	If there aren't any active users using the Launchpad application, the system stops the pod after the set number of minutes.	N/A

Application	Property	Description	Default Setting	Impact	More Information
ML Tracking	Mounted tls certificate path	Specifies the location of your SAP HANA certificates for TLS protocol.	/etc/certs/root-ca/ca-bundle.pem	N/A	N/A
ML Tracking	Use TLS	Specifies whether to connect to SAP HANA database, which stores the metrics, using TLS encryption.	True (on)	N/A	N/A
ML-API	Limit for the maximum size of a file in the workspace of an ML Scenario (bytes)	Specifies the maximum size allowed for files in the workspace of an ML Scenario.	2097152 bytes	N/A	N/A
ML-API	Connection timeout for outgoing calls made from ML Scenario API (seconds)	Specifies the number of seconds that the application waits for a response after connecting to an external service before returning an error.	10 seconds	N/A	N/A
ML-API	Maximum amount of time for ML Scenario API to respond to incoming calls, before being declared unhealthy (seconds)	Specifies the number of seconds that the application waits for an internal signal before declaring itself as unresponsive.	600 seconds	N/A	N/A
ML-API	Read timeout for outgoing calls made from ML Scenario API (seconds)	Specifies the number of seconds that the application waits for a response after making a request to an external service before returning an error.	30 seconds	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
ML-API	Logging level for external components logs within ML Scenario API	<p>Indicates the type of logging level for external components logs within ML Scenario application.</p> <p>Possible values include the following:</p> <ul style="list-style-type: none"> • Critical • Error • Warning • Info • Debug <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p>	Warning	N/A	N/A
ML-API	Logging level for ML Scenario API logs	<p>Indicates the type of logging level for ML Scenario application.</p> <p>Possible values include the following:</p> <ul style="list-style-type: none"> • Critidal • Error • Warning • Info • Debug <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p>	Info	Setting this configuration to Debug can lead to performance degradation, as the amount of logs may increase significantly.	N/A

Application	Property	Description	Default Setting	Impact	More Information
ML-API	Read timeout for outgoing calls made from ML Scenario API to Pipeline Modeler (seconds)	Specifies the number of seconds that the application waits for a response from the Modeler after contacting it before returning an error.	30 seconds	N/A	N/A
ML-API	Logging level for SQL-related logs within ML Scenario API	<p>Indicates the type of logging level for SQL-related logs within ML Scenario application.</p> <p>Possible values include the following:</p> <ul style="list-style-type: none"> • Critical • Error • Warning • Info • Debug <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p>	Warning	N/A	N/A
ML-API	Timeout for starting Pipeline Modeler executions from ML Scenario API (seconds)	Specifies the number of seconds that the application waits for a response from the Modeler after contacting it before returning an error.	100 seconds	N/A	N/A
System Management	Enable http access log	Specifies whether to log HTTP request and responses made to the System Management application.	True (on)	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
System Management	http client request max body size	Maximum size of the HTTP request body and maximum file size that can be uploaded to the System Management application.	1M	N/A	N/A
Modeler	CPU request for the vflow apps, modification only advised in small clusters	<p>Specifies the default CPU requests for the following vFlow cluster applications:</p> <ul style="list-style-type: none"> vflow vflow-ctrl-main vflow-ctrl-sub vflow-runtime-store vflow-image-build vflow-diagnostics <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	50 millicores	In small clusters, to allocate more resources for running graphs, you can reduce this value.	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Average CPU usage to trigger the pod autoscaler	<p>Specifies the average utilization of CPU for the following vFlow pods:</p> <ul style="list-style-type: none"> vflow vflow-runtime-store vflow-image-build vflow-diagnostics <p>These vFlow pods activate Kubernetes horizontal pod autoscaling (for more information, see horizontal pod autoscaling).</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	400 millicores	N/A	N/A
Modeler	AWS IAM role for package version	<p>Provides necessary permission to the Modeler to access AWS resources when building and uploading graph Docker images.</p> <p>Deprecated in SAP Data Intelligence 2022.21.2.</p>	N/A	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Docker registry for modeler base images	Cluster-level parameter that can be changed only by the cluster administrator upon request. Deprecated in SAP Data Intelligence version DI:2023.03:E.	N/A	N/A	N/A
Modeler	Garbage collection interval for build logs	Specifies the elapsed time between which the system removes old or unnecessary build logs. Cluster-level parameter that can be changed only by the cluster administrator upon request. Deprecated in SAP Data Intelligence version 2023.03.	2 hours	Frees space and keeps the system running smoothly. Increasing or decreasing this value can affect the performance and stability of the system.	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Default resources definition for kaniko pod	<p>Specifies the CPU and memory resources allocated to a kaniko pod.</p> <p>Kaniko is a tool that builds Docker images in a Kubernetes cluster without having to use a Docker daemon. Each Docker image can have its own custom-built resources defined inside 'resourceDefinition.json', which is saved in the same folder as Dockerfile.</p>	<pre>{ "requests" : { "cpu" : "1.0" , "memory" : "2Gi" } }</pre>	Custom resources override the Default resources definition for kaniko pod .	N/A
Modeler	Default Timeout for calling external service to get diagnostics file	<p>Specifies the maximum time that the system waits for a response from the external service to get a diagnostics file before considering the request as failed.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence version 2023.03.</p>	60 seconds	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Garbage collection time limit for finished graphs	Specifies the time before the garbage collector frees memory used by the finished graphs.	72 hours	Increasing the value can affect the performance and stability of the system. If you use number-based garbage collection strategy, the Modeler removes finished graphs when either or both conditions are fulfilled.	For more information about garbage collection, see Graph Execution Garbage Collection in the <i>Modeling Guide</i> .
Modeler	Enable garbage collection for dead graphs	Specifies whether the Modeler removes dead graphs from the system.	True (on)	Changing this parameter to False (off) can affect the performance and stability of the system.	N/A
Modeler	Enable garbage collection for paused graphs	Specifies whether the Modeler garbage collector removes paused graphs from the system.	False (off)	N/A	For more information about garbage collection, see Graph Execution Garbage Collection in the <i>Modeling Guide</i> .
Modeler	Garbage collection for snapshot graphs	Specifies whether the Modeler removes graphs with snapshots enabled from the system when the graph completes. Cluster-level parameter that can be changed only by the cluster administrator upon request.	True (on)	Setting the value to False can affect the performance and stability of the system.	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Maximum number of graphs per user. Finished graphs will be deleted when exceeded	Specifies a limit for the number of graphs per user in the Modeler at any time. When a user exceeds the limit, the Modeler removes completed graphs.	50 graphs	<p>Increasing the value can affect the performance and stability of the system.</p> <p>Maximum allowed graphs is 500. You should regularly purge old dead graphs that you no longer need to free up space for new graphs to persist. After the limit is reached, any new graphs are immediately garbage collected upon completion and their lineage extraction fails with an error.</p>	For more information about garbage collection, see Graph Execution Garbage Collection in the <i>Modeling Guide</i> .
Modeler	Maximum number of concurrency hana connections for graph gc.	<p>Specifies the maximum number of concurrent connections to the internal SAP HANA database that the system allows for a given graph.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	20 connections	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Graph execution history retention time limit	<p>Specifies the maximum time that the Modeler retains the execution history of a graph.</p> <p>History is a subset of data from the graph that is stored after it's deleted.</p>	2160 hours	Setting a retention time limit ensures that the execution history doesn't become too large. Increasing the value can affect the performance and stability of the system.	N/A
Modeler	Timeout for a pod to be scheduled	<p>Specifies the time that the system waits for graph pod or pods to become available before returning an error.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	3 minutes	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Timeout for a[n] image build to be in running state	<p>Specifies the maximum time the Modeler waits for a response from the image-build service before considering the request as failed.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 5px; margin: 10px 0;"> <p>Note</p> <p>The image-build service builds graph Docker images.</p> </div> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	12 hours	N/A	N/A
Modeler	Disable image check	<p>Deprecated.</p>	N/A	N/A	N/A
Modeler	Docker image pull secret for Modeler	<p>Allows the Modeler to pull graph images from authenticate Docker repository.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	N/A	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Internal cluster settings, change only when advised by SAP product support	<p>Various cluster parameters that aren't meant to be changed by non administrator users.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	N/A	N/A	N/A
Modeler	Internal tenant settings, change only when advised by SAP product support	<p>Allows setting of various tenant-specific parameters, such as the following:</p> <ul style="list-style-type: none"> • Default resources for group. • Custom bus image. • vFlow build image. <p>Protected parameter that can be changed only by the cluster administrator upon request.</p>	List of key-value pairs: parameter : value	N/A	N/A
Modeler	Kaniko build container image	Protected parameter.	N/A	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Maximum request body size in megabytes	<p>Specifies the maximum body size sent in a single request.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	10 Mb	N/A	N/A
Modeler	Wait for last operator metrics to be collected before shutdown of group	Specifies to wait for collection of metrics from the last operator before shutdown of the group.	True (on)	<p>Waiting ensures that all necessary data is available for analysis and troubleshooting, if necessary.</p> <p>Turning off this parameter (false) can make the graph termination on the nodes faster, but guarantees on the graph metrics during shutdown are lost.</p>	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Enable Prometheus metric publisher for Modeler	<p>Allows the Modeler to collect metrics about graph performance and behavior and send to a Prometheus server for storage and analysis.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	True (on)	N/A	N/A
Modeler	Docker registry for Modeler images	<p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	N/A	N/A	N/A
Modeler	Name of the vSystem secret containing the credentials for Docker registry	<p>Cluster-level parameter that can be changed only by the cluster administrator.</p>	N/A	N/A	N/A

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Trace level	<p>Indicates the type of logging level for vFlow apps. Possible values include:</p> <ul style="list-style-type: none"> • None • Fatal • Error • Warning • Info • Debug <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	Info	N/A	For descriptions of all trace levels, see Trace Severity Levels .
Modeler	Timeout to be used for vSystem client API calls	<p>Specifies the maximum time that the Modeler waits for a response from any vSystem service before considering the request as failed.</p> <p>Cluster-level parameter that can be changed only by the cluster administrator upon request.</p> <p>Deprecated in SAP Data Intelligence 2023.03.</p>	3 minutes	N/A	N/A
Modeler	Enable save diagnostics archive to external storage defined via Connection ID	<p>Saves diagnostics archive for graphs in external storage.</p>	False (off)	N/A	Saving Diagnostic Information for Graphs on External Storage

Application	Property	Description	Default Setting	Impact	More Information
Modeler	Connection ID for diagnostics storage	Specifies the connection ID for storage that contains the diagnostic archive.	DI_DATA_LAKE	N/A	Saving Diagnostic Information for Graphs on External Storage
Modeler	Store diagnostics in external storage only for failed graphs	Specifies to use only external storage for the diagnostic archive of failed graphs.	True (on)	N/A	Saving Diagnostic Information for Graphs on External Storage
Modeler	Garbage collector strategy for diagnostics storage	Specifies the strategy for removing the diagnostics archive from storage. Options are: never or history. History refers to the moment that a historical graph is removed from the system, which is determined by the execution history retention time limit parameter.	history	N/A	Saving Diagnostic Information for Graphs on External Storage
System	App base Docker image	This property is view only. You can't change the value.	N/A	N/A	N/A
System	App base image pull secret	This property is view only. You can't change the value.	N/A	N/A	N/A
System	HTTP proxy	This property is view only. You can't change the value.	N/A	N/A	N/A
System	HTTPS proxy	This property is view only. You can't change the value.	N/A	N/A	N/A
System	Pull secret for apps images	This property is view only. You can't change the value.	N/A	N/A	N/A
System	Kubernetes namespace	This property is view only. You can't change the value.	N/A	N/A	N/A






Application	Property	Description	Default Setting	Impact	More Information
System	No proxy	This property is view only. You can't change the value.	N/A	N/A	N/A
System	Docker registry with apps images	This property is view only. You can't change the value.	N/A	N/A	N/A
System	Version	This property is view only. You can't change the value.	The currently installed System Management version.	N/A	N/A
System	SAP Data Intelligence version	This property is view only. You can't change the value.	The currently installed SAP Data Intelligence version.	N/A	N/A
System	Enable password login	Specifies whether password login is enabled for a tenant.	True	If set to false, password login is disabled for the tenant, and login is possible only by other authentication methods such as via a certificate.	N/A
VsystemUI	Node V8 Options (bytes)	Specifies the maximum size of HTTP headers. Can be set only by system tenants.	16000	The total size of HTTP headers received by Node.js server must not exceed the specified number of bytes.	N/A

4 Using SAP Data Intelligence Connection Management

SAP Data Intelligence administrators or other business users with necessary privileges can use the SAP Data Intelligence Connection Management application to create and maintain connections. A connection represents an access point to a remote system or a remote data source.

The table lists the various actions that you can perform in the Connection Management application.

Actions	Description
Create Connections	<ol style="list-style-type: none">1. Launch the Connection Management application.2. Open the <i>Connections</i> tab.3. Select + (<i>Create Connection</i>). <p>For more information, see Create a Connection [page 120].</p>
Filter Connections	<p>Filter connections based on tags and connection types. A tag is an attribute to group and filter connections. Each connection type is associated with a fixed set of predefined tags (storage, application, http, and db).</p> <ol style="list-style-type: none">1. Open the <i>Connections</i> tab.2. Select ∇ (<i>Filter</i>).3. Define the filter conditions by choosing tags and types.4. Select <i>OK</i>.
Edit a Connection	<p>To edit an existing connection, follow these steps:</p> <ol style="list-style-type: none">1. Open the <i>Connections</i> tab, and choose the required connection.2. Select ⋮ (<i>Select an Action</i>) in the Actions column, and select ✎ (<i>Edit</i>).3. Edit the connection details as necessary.4. Select <i>Save</i>. <div data-bbox="730 1541 1396 1729" style="border-left: 2px solid #0070C0; padding-left: 10px;"><p>Note</p><p>You can't change the connection ID or the connection type. If you don't have sufficient permissions to edit a connection, the <i>Save</i> button isn't available.</p></div>
Delete a Connection	<p>To delete an existing connection, follow these steps:</p> <ol style="list-style-type: none">1. Open the <i>Connections</i> tab, and choose the required connection.2. Select ⋮ (<i>Select an Action</i>) in the Actions column, and select 🗑 (<i>Delete</i>).3. Select <i>OK</i> to confirm the delete operation.

Actions	Description
View Connection Status	<p>To view a connection status, follow these steps:</p> <ol style="list-style-type: none"> 1. Open the <i>Connections</i> tab, and choose the required connection. 2. Select  (<i>Select an Action</i>) in the <i>Actions</i> column, and select  (<i>Check Status</i>). <p>The status check performs a type-specific reachability test of the remote system referred to in the connection definition. The possible connection statuses are OK, ERROR, and UNKNOWN. The application displays the status in the <i>Connection Status</i> dialog box.</p>
Import or export connections	<p>Create a connection in SAP Data Intelligence by exporting or importing a connection as a JSON schema.</p> <p>To export one or more connections, follow these steps:</p> <ol style="list-style-type: none"> 1. Open the <i>Connections</i> tab. 2. Select  (Export). 3. Choose the connection ID or IDs in the <i>Export Connections</i> dialog box. 4. Select <i>Export</i>. <p>Connection Management exports the selected connection JSON schema or schemas to your local system.</p> <p>To import a connection:</p> <ol style="list-style-type: none"> 1. Open the <i>Connections</i> tab. 2. Select  (Import). 3. Browse for and select the JSON schema of the connection to import. <div data-bbox="730 1339 1394 1523" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p> Note</p> <p>The username and password details aren't imported. You have to manually provide this information after importing and create a connection.</p> </div>

Related Information

[Log on to SAP Data Intelligence Connection Management \[page 120\]](#)

[Create a Connection \[page 120\]](#)

[Manage Certificates \[page 122\]](#)

[Supported Connections in SAP Data Intelligence \[page 205\]](#)

[Using SAP Cloud Connector Gateway \[page 223\]](#)

[\(Mandatory\) Configure Authorizations for Supported Connection Types \[page 224\]](#)

[Allowing SAP Data Intelligence Access Through Firewalls \[page 225\]](#)

4.1 Log on to SAP Data Intelligence Connection Management

Access the SAP Data Intelligence Connection Management application after you log on to SAP Data Intelligence.

Procedure

1. Enter the SAP Data Intelligence URL in a browser.
2. Enter your login credentials.

The SAP Data Intelligence Launchpad opens. The home page displays the applications available in the tenant based on the policies assigned to you.

Note

If you enter an incorrect password five consecutive times within a minute, your account will be temporarily locked for 10 seconds until your next attempt.

3. Select the *Connection Management* tile.

The Connection Management application opens displaying the initial screen.

4.2 Create a Connection

Create a connection in SAP Data Intelligence, which represents an access point to a remote system or a remote data source.

Context

Any user with member privileges can create, view, edit, and delete their own connections. The connections are private, which means they aren't visible to other member users, but they're viewable by the tenant administrator. To grant access to another user, the tenant administrator must assign a corresponding policy to that user.

Create connections in the SAP Data Intelligence Connection Management application. The application associates each connection with a unique ID and a type that specifies the nature of the access point. In addition to the required unique ID attribute, every connection can have an optional description. After creating a connection, use the connection ID in the SAP Data Intelligence Modeler to reference an external resource.

To create a connection in SAP Data Intelligence, perform the following steps:

Procedure

1. Start the SAP Data Intelligence Connection Management application.
2. Choose *Create* in the editor toolbar.

The *Create Connection* page opens.

3. Enter a unique connection ID in the *ID* text box.

The connection ID can contain uppercase letters, lowercase letters, digits, or underscores.

4. Select the applicable connection type from the *Connection Type* list.

The connection type determines the rest of the connection information that you provide.

Note

For more information about connection types, see the topic [Supported Connection Types \[page 122\]](#).

Tenant administrators or graph modelers can troubleshoot connectivity issues when they know the user name and database credentials for the connection. To view the user name of the connection, and to obtain the credentials for accessing the connected data sources, enable the `app.datahub-app-core.connectionCredentialsUnmasked` policy. This policy provides full visibility to selected connection fields, such as user names.

5. Enter the remaining connection information.
6. Add custom tags by performing the following substeps:
Associate a connection with one or more user-defined tags, which can help you with filter connections, and so on. After creating a tag, you can reuse it in other connection definitions.

Example

Associate the connection with tags, such as **db**, **storage**, and **application**. Later, you can use these tags to filter connections in the application user interface.

- a. Enter a tag in the *Tag* text field, or select an existing tag from the list.
 - b. Press to define or reuse another tag.
7. Select *Test Connection* to test the connection.

If the application can create a connection to the target system successfully, it displays the connection status as *OK*. Otherwise, it displays an error message. If there are errors, verify the connection details that you've provided and test again.

Restriction

The *Test Connection* option isn't available for HTTP connection type.

8. Select *Create*.
9. Select *Edit* to change the connection settings, if necessary.

4.3 Manage Certificates

Use the SAP Data Intelligence Connection Management application to manage certificates for remote systems.

Context

The connection management application provides capabilities to import certificates for connections to remote systems. For operators in the SAP Data Intelligence Modeling tool that leverage HTTPS as the underlying transport protocol (using TLS transport encryption), the certificate of the upstream system must be trusted. To import a certificate into the trust chain, obtain the certificates from the target endpoint and import it using the connection management application.


Procedure

1. Start the SAP Data Intelligence Connection Management application.
2. On the home page, choose the [Certificates](#) tab.
3. In the menu bar, choose [Import](#) icon to import a certificate.
4. Browse and select your certificate.

The application displays the list of certificates imported under the [Certificates](#) tab.

Note

Certificates uploaded into the SAP Data Intelligence Connection Management application are not imported to the `ca` folder in the SAP Data Intelligence System Management application under [Files Workspaces](#).

5. **Optional:** To delete a certificate, hover on the certificate and choose .

Remember

You can't delete the preinstalled certificates.

4.4 Supported Connection Types

SAP Data Intelligence delivers a set of predefined connection types. These connection types represent a specific category of remote resource.

The table lists the supported connection types in SAP Data Intelligence.

Supported Connection Types

Connection Type	Connection
ABAP [page 124]	ABAP
ABAP LEGACY [page 127]	ABAP Legacy
ADL (Deprecated) [page 131]	Microsoft Azure Data Lake
ADL_V2 [page 132]	Microsoft Azure Storage Gen2
AWS_SNS [page 134]	Amazon Simple Notification Service
AZURE_SQL_DB [page 135]	Microsoft Azure Cloud SQL Database
BW [page 137]	SAP Business Warehouse
CLOUD_DATA_INTEGRATION [page 138]	OData-based APIs for data integration with SAP Data Intelligence and SAP BW/4 HANA
CPEM [page 140]	SAP Event Mesh (a capability of SAP Integration Suite)
CPI [page 141]	HTTPS connection to SAP Integration service that runs on SAP Integration Suite
DATASERVICES [page 142]	SOAP server and information from an SAP Data Service Administration server
DB2 [page 143]	IBM DB2 databases
GCP_BIGQUERY [page 145]	Google Cloud BigQuery
GCP_DATAPROC [page 148]	Google Cloud Dataproc cluster
GCP_PUBSUB [page 149]	Google Cloud Platform publish/subscribe service
GCS [page 149]	Google Cloud Storage
HANA_DB [page 150]	SAP HANA database
HANA_XS [page 153]	Tables and views in SAP HANA databases.
HDFS [page 154]	Hadoop Distributed File System server
HDL_DB [page 156]	SAP HANA Data Lake database
HDL_FILES [page 159]	SAP HANA Data Lake information
HTTP [page 160]	Servers over HTTP or HTTPS
IMAP [page 161]	Internet Message Access Protocol server
INFORMATION_STEWARD [page 162]	SAP Information Steward Administration server
KAFKA [page 163]	Apache Kafka cluster
MSSQL [page 169]	Microsoft SQL Server database
MYSQL [page 170]	Oracle MySQL server
ODATA [page 171]	OData RESTful API
OPEN_CONNECTORS [page 173]	Open Connectors, a core capability of SAP Integration Suite
OPENAPI [page 175]	OpenAPI 2.0 specification server
ORACLE [page 176]	Oracle database
OSS [page 179]	Alibaba Cloud Object Storage Service

Connection Type	Connection
POSTGRESQL [page 181]	PostgreSQL database
REDSHIFT [page 182]	Amazon Redshift databases
RSERVE [page 186]	RServe server
S3 [page 186]	Amazon Simple Storage Service
SAP_IQ [page 189]	SAP IQ databases
SDL [page 192]	Remote object stores
SFTP [page 194]	SSH File Transfer Protocol server
SMTP [page 196]	Simple Mail Transfer Protocol server
SNOWFLAKE [page 197]	Snowflake databases
TERADATA [page 201]	Teradata databases

Related Information

[Changing Data Capture \(CDC\)](#)

4.4.1 ABAP

An ABAP connection type connects to and accesses information from objects (tables, CDS views, and ODP objects) in an ABAP-based SAP system (SAP S/4HANA, SAP S/4HANA Cloud, or SAP ECC).

Note

When creating an ABAP connection using the SAP Data Intelligence Connection Management application, save the connection before you perform a connection check. The application runs the check based on the connection information that you save. If the connection check fails, double-check your entries to make sure that you have valid entries in all relevant fields for your use case.

For more information about the ABAP connection type, see SAP Note [2835207](#) .

Operations

The ABAP connection type allows for the following operations:

- Access data from ABAP-based SAP systems.
- Access metadata of supported objects.
- Work with operators in SAP Data Intelligence Modeler based on the ABAP language.

Attributes

Attribute	Description
Protocol	Protocol type for connecting to the ABAP system. Options include the following: <ul style="list-style-type: none">• RFC• WebSocket RFC
Connection Type	Connection type for ABAP connection. Options include the following: <ul style="list-style-type: none">• With Load Balancing: Uses the ABAP connection against the message server of the ABAP source system.• Without Load Balancing: Uses the ABAP connection against a specific application server of the ABAP source system.
System ID	ABAP system ID. For more information about the system ID, see your ABAP documentation.
Application Server	Hostname or IP address of the application server. Applicable when you set Protocol to RFC.
Instance Number	Two-digit instance number for the RFC connection. Applicable when you set Protocol to RFC.
Client	Client of the ABAP system to use for logon.
Authentication	Authentication type to use. Authentication options include the following: <ul style="list-style-type: none">• Basic: Select when you use a standard RFC connection to connect to the ABAP system. Also specify values for User and Password to access the ABAP system.• ClientCertificate: Select when you use an RFC connection that uses the WebSocket protocol to connect to the ABAP system. Also complete the Client Certificate attribute.
Enable SNC	Specify whether to use Secure Network Communication (SNC) to establish RFC connection. Applicable when you set Protocol to RFC.
User	User name required for authentication. Applicable when you set Authentication to Basic.
Password	Password that corresponds to the specified user name for authentication. Applicable when you set Authentication to Basic.

Attribute	Description
Language	Two-digit ISO language code. If you leave this attribute blank, the application uses the default logon language of the ABAP system.
Gateway Host	Host name of the Gateway server for the RFC connection. Applicable when you set Protocol to RFC.
SAP Router	SAP router configuration string for the RFC connection. The application ignores this value when you use the Cloud Connector. Applicable when you set Protocol to RFC.
Hostname	Host name of the WebSocket RFC connection endpoint. Applicable when you set Protocol to WebSocket RFC.
Portnumber	Port number of the WebSocket RFC connection endpoint value. The default port number is 443. Applicable when you set Protocol to WebSocket RFC.
Client Certificate	X.509 client certificate to access the ABAP system. Applicable when you set Protocol to WeSocket RFC and Authentication to ClientCertificate.
Client Private Key	Applicable when you set Protocol to WeSocket RFC and Authentication to ClientCertificate.
Client Private Key Password	Applicable when you set Protocol to WeSocket RFC and Authentication to ClientCertificate.
SNC Partner Name	SNC name of the communication partner. Applicable when you set Enable SNC to true (on).
Quality of Protection	SNC protection level: <ul style="list-style-type: none"> • 1: Apply authentication only. • 2: Apply integrity protection - includes level 1, authentication. • 3: Apply privacy protection - includes level 2, integrity protection and level 1, authentication. • 8: Apply the default protection. • 9: Apply the maximum protection. Applicable when you set Enable SNC to true (on).
Message Server Host	Host name or IP address of the message server for RFC connection.
Message Server Port	Port number of the message server for RFC connection.
Logon Group	Default value is SPACE.

Connection information is persisted in the connection cache, therefore, changes to a connection doesn't take effect immediately after the change.

Note

For information about supported versions, see [Installation-Related Information](#).

4.4.1.1 Configure SNC for ABAP

There are two methods to connect an on-premise ABAP-based SAP system to SAP Data Intelligence Cloud using SNC (Secure Network Communication).

Configure Cloud Connector

Use the Cloud Connector to connect an on-premise ABAP-based SAP system to SAP Data Intelligence Cloud using SNC. First, configure SNC for Cloud Connector as described in the topics [Initial Configuration \(RFC\)](#) and [Configure Principal Propagation for RFC](#).

Make sure that you set the connection attribute *Enable SNC* to **off (false)** in the ABAP connector type configuration in SAP Data Intelligence Cloud. Deactivating SNC ensures that SNC is consistently taken care of by the Cloud Connector in the connected SAP system.

X.509 Certificate

Use an existing X.509 certificate to connect an on-premise ABAP-based SAP system to SAP Data Intelligence Cloud using SNC. Configure the ABAP connection type in the SAP Data Intelligence Cloud Connection Management and make sure to include the following settings:

- Choose **ClientCertificate** for *Authentication*.
- Set *Enable SNC* to **on (true)** and enter values for *Client Certificate*, *Private Key*, and *Private Key Password*.

4.4.2 ABAP LEGACY

The ABAP LEGACY connection type connects to and accesses information from ODP objects in an SAP ABAP system where DMIS installation isn't possible.

Note

ABAP LEGACY doesn't provide all the features that an ABAP connection provides:

- Supports only the RFC protocol.
- Supports Password and Secure Network Communications(SNC) authentication.

- Doesn't support metadata extraction and lineage through Metadata Explorer.
- Browses and previews only ODP objects.
- Doesn't support ABAP TABLES.
- Supports only the *SAP Application Consumer* operator and the *Flowagent ABAP ODP Consumer* operator.

For more information about ABAP connection types, see SAP Note [2835207](#).

On-Premise Connectivity Through Cloud Connector

For on-premise connectivity through the Cloud Connector (for SAP Data Intelligence cloud installation), either add the prefixes to the allowlist (/SAPDS/, BAPI, RODPS and RFC), or add the exact functions to the allowlist.

Select [Cloud Connector Administration](#) > [Cloud To On-Premise](#) > [Access Control](#).

The following are the exact functions:

- /SAPDS/EXTRACTOR_NAVIGATE
- /SAPDS/GET_VERSION
- /SAPDS/MODEL_NAVIGATE
- BAPI_USER_GET_DETAIL
- RFC_FUNCTION_SEARCH
- RODPS_REPL_CONTEXT_GET_LIST
- RODPS_REPL_ODP_OPEN
- RODPS_REPL_ODP_FETCH
- RODPS_REPL_ODP_GET_DETAIL
- RODPS_REPL_ODP_GET_LIST

SNC Support

To configure SNC support, see [Configure SNC for ABAP_LEGACY \[page 130\]](#).

SNC support limitations:

- You can't use authentication type "SNC" with Gateway-enabled ABAP_LEGACY connection.
- You can't configure SNC on Cloud Connector with Username and Password authentication type on the ABAP_LEGACY connection.

Operations

The ABAP LEGACY connection allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- View fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Run rules in Metadata Explorer.
- Read tables and views in Modeler.

Attributes

Attribute	Description
Protocol	RFC is the only protocol option for ABAP LEGACY. For other types, use the ABAP connection.
System ID	For more information about ABAP LEGACY system ID, see your ABAP documentation.
Application Server	Host name or IP address of the application server for RFC connection.
Client	Client of the ABAP system to use for logon.
Instance Number	Two-digit instance number.
Gateway Host	Host name of the Gateway server for RFC connection.
SAP Router	SAP router configuration string for RFC connection.
Authentication	Authentication types: <ul style="list-style-type: none"> • Password • SNC
User	User name that the application uses for authentication. Applicable when you select Authentication = Password.
Password	Password that the application uses for authentication. Applicable when you select Authentication = Password.
SNC Client Name	Name of the secure network communication client. Applicable when you select Authentication = SNC.
SNC Partner Name	Name of the secure network communication partner. Applicable when you select Authentication = SNC.
Client Certificate	Client certificate with key in PEM. Applicable when you select Authentication = SNC.
ABAP Server Certificate	Applicable when you select Authentication = SNC.
Language	Two-digit ISO language code. If not set, the application uses the default logon language of the ABAP system.

Attribute	Description
Connection Type	Indicates whether the connection has load balancing: <ul style="list-style-type: none"> • With Load Balancing • Without Load Balancing
Message Server Host	Host name or IP address of the message server for the RFC connection.
Message Server Port	Port number of the message server for the RFC connection.
Logon Group	The default value is SPACE.

4.4.2.1 Configure SNC for ABAP_LEGACY

When you choose to use SNC (Secure Network Connection) authentication for an ABAP_LEGACY connection type, you must first enable SNC before you create the ABAP_LEGACY connection.

Context

To enable SNC for ABAP_LEGACY connection type for one user using self-signed, perform the following steps:

Procedure

1. Create self-signed or CA-signed client certificate for a user.

Example

For a user named DI_USER, create a self-signed certificate with a PEM file that contains the key and certificate, and then use the client PEM file DI_USER.pem (with key) for the ABAP_LEGACY connection type:

Sample Code

```
openssl req -x509 -sha256 -nodes -days 1825 -newkey rsa:2048 -keyout
DI_USER.key -out DI_USER.crt -subj "/C=US/O=YOURORG/CN=DI_USER"
//export client certificate, key to pfx
openssl pkcs12 -export -out DI_USER.pfx -inkey DI_USER.key -in
DI_USER.crt
//convert pfx to pem
openssl pkcs12 -in DI_USER.pfx -out DI_USER.pem -nodes
```

The application imports the client certificate file DI_USER.crt (contains only the public key) to the ABAP system using /STRUST transaction.

2. Select *Import Certificate* and choose the certificate file.
3. Export the server certificate from ABAP server as base64, and save as `SERVER.crt`.

The application uses the server certificate as a server in the ABAP_LEGACY connection. The application uses the subject name of the server for the SNC partner name in the SAP Data Intelligence connection.

❖ Example

`p:CN=GCX, O=SAP-AG, C=DE`

It's important to include the spaces in between key pair values.

4. Update the `DI_USER` SNC name by using transaction `/su01`.
Transaction `/su01` is the SNC Name (client) for the SAP Data Intelligence connection. The SNC name for the SAP Data Intelligence connection is `p:CN=DI_USER, O=YOURORG, C=US`. You can take this value from `STRUST` (Trust manager) when you import the client certificate.

Results

The ABAP server-side setup is complete, and all the required details for the ABAP_LEGACY connection is available.

4.4.3 ADL (Deprecated)

The ADL connection type connects to and accesses information from objects in Microsoft Azure Data Lake (ADL).

⚠ Caution

Microsoft Azure Data Lake Storage Gen1 will be retired on February 29, 2024. After this date, it might not be possible to access data with Gen1 connections. Therefore, the connection type is deprecated and may be removed in future SAP Data Intelligence releases.

After your Microsoft Azure Data Lake Storage Gen1 has been migrated to Gen2, recreate the connection with the Gen2 connection type (`ADL_V2`) and update your graphs and operators to the new connection.

For more information, see [Action required: Switch to Azure Data Lake Storage Gen2 by 29 February 2024](#) on the Azure web site.

Operations

The ADL connection type allows for the following operations:

- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.

- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Save data from the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Read and write files using the Read File and Write File operators in SAP Data Intelligence Modeler.
- Rename and remove files using the Move File and Remove File operators in Modeler.

Attributes

Attribute	Description
Account name	Name of the Azure Data Lake Storage Gen1 account.
Tenant ID	ID of the Azure Data Lake Storage Gen1 tenant.
Client ID	Client ID, also referred to as Application ID.
Client Key	Client key, also referred to as Client Secret or Authentication Key.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.4 ADL_V2

The ADL_V2 connection type connects to and accesses information from objects in Microsoft Azure Storage Gen2.

Operations

The ADL_V2 connection type allows for the following operations:

- Browse folders and files in the Azure Data Lake Storage Gen2.
- Obtain file metadata.
- Profile data.
- Preview data.
- Perform flowgraph tasks with Azure Data Lake Storage Gen2 files as the source and target.
- Read and write files in byte chunks with the Read File and Write File operators in SAP Data Intelligence Modeler.

- Read and write CSV, Parquet, and ORC files using Structured File Consumer and Structured File Producer operators in Modeler.
- Rename and remove files with the Move File and Remove File operators in Modeler.

For more information about Azure credentials, see the following Microsoft documentation: <https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal#get-values-for-signing-in> .

Attributes

Attribute	Description
Authorization Method	Authorization method for this connection. Options include the following: <ul style="list-style-type: none"> • shared_key • shared_access_signature • oauth2_client_credentials • oauth2_username_password
Account Key	Account key required when you choose the shared_key Authorization Method.
Account Name	Account name required when you choose the shared_key Authorization Method.
Endpoint Suffix	Endpoint suffix applicable for any of the listed Authorization Methods. If you leave this blank, the application uses the default value <code>core.windows.net</code> .
Root Path	Root path for browsing. Start the path with a forward slash character followed by the file system name. For example, <code>/MyFileSystem/MyFolder</code> . The application prefixes this root path to any path with this connection.
SAS Token	Required when you choose the shared_access_signature Authorization Method.
Use Client Certificate	Specifies whether to use client certificate for authorization. The default is false (off). Applicable when Authorization Method is <code>oauth2_client_credentials</code> .
Oauth2TokenEndpoint	Token endpoint to get the access token. Applicable when Authorization Method is <code>oauth2_client_credentials</code> .
Oauth2ClientId	Client ID used in the Client Credentials with Client Secret and Client Certificate authentication. Applicable when Authorization Method is <code>oauth2_client_credentials</code> .
Client Certificate	X.509 client certificate used in the Client Certificate authentication. Applicable when Authorization Method is <code>oauth2_client_credentials</code> .

Attribute	Description
Client Private Key	X.509 client's private key used in the Client Certificate authentication. Applicable when Authorization Method is <code>oauth2_client_credentials</code> .
Oauth2ClientSecret	Client secret used in the Client Credentials with Client Secret authentication. Applicable when Authorization Method is <code>oauth2_client_credentials</code> .
Oauth2ClientEndpoint	Client endpoint to get the access token for the Authorization Method <code>oauth2_username_password</code> .
Oauth2Username	User name for the Username and Password authentication method. Applicable when Authorization Method is <code>oauth2_username_password</code> .
Oauth2Password	Password for the Username and Password authentication method. Applicable when Authorization Method is <code>oauth2_username_password</code> .

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.5 AWS_SNS

The `AWS_SNS` connection type connects to and accesses information from the Amazon Simple Notification Service (SNS). Amazon SNS is a managed pub/sub service.

Operations

The `AWS_SNS` connection type allows for the following operations:

- Create a topic in the AWS SNS service.
- Subscribe to a topic on the AWS SNS service.
- Send messages to topics using the AWS SNS Producer operator in SAP Data Intelligence Modeler.
- Receive messages from the topics using the AWS SNS Consumer operator in Modeler.

Attributes

Attribute	Description
AWS Account ID	Account ID assigned to the IAM user that owns the SNS resources.

Attribute	Description
AWS Access Key	ID of the key that accesses the AWS SNS API. The application uses this ID in combination with the secret key. The key must have permissions to use the SNS service.
AWS Secret Key	Key that accesses the AWS SNS API. The application uses this key in combination with the key ID. The key must have permissions to use the SNS service.
Region	Region where the application stores or looks up SNS topics and subscriptions.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.6 AZURE_SQL_DB

The AURE_SQL_DB connection type connects to and accesses information from a Microsoft Azure Cloud SQL Database using SQL Server authentication.

Prerequisites

The network for Azure SQL DB instances is protected by a firewall that controls incoming traffic.

SAP Data Intelligence Connection Management exposes an IP address through a read-only connection called `INFO_NAT_GATEWAY_IP`. Use this IP address and add the connection to the allowlist in the Azure dashboard. To add domain IPs to the allowlist, perform the following steps:

1. Access your Azure SQL Server.
2. Choose the SQL databases in **Settings**, and then choose the database with which to connect.
3. Select **Set server firewall** and then select **+ (Add client IP)**.
4. Specify a rule name and the IP range of the Kubernetes nodes.

Operations

The AZURE_SQL_DB connection allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.

- Prepare data using the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in SAP Data Intelligence Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Host	Fully qualified server name or IP address for the Azure SQL database.
Port	Port number of the Azure SQL database server.
Validate host certificate	Specify whether to validate the server certificate in your certificate library.
Hostname in certificate	Host name on the host certificate. The application verifies the specified host name against the host name in the certificate.
Database name	Name of the Azure_SQL_DB with which to connect.
User	User ID of the SQL login with SQL authentication.
Password	Password from the server administrator account that was used to create the server.
Additional session parameters	Additional session variables, if necessary.

Note

SQL authentication is the only method supported by SAP.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.7 BW

The BW (SAP Business Warehouse) connection type connects to and accesses information from SAP BW systems, such as SAP BW, SAP BW on HANA and SAP BW/4 HANA.

Operations

The BW connection type allows for the following operations:

- Browse and preview InfoProviders and BW Queries as datasets in Metadata Explorer.
- Run BW process chains using the BW Process Chain operator in the SAP Data Intelligence Modeler.
- Move data from InfoProviders and BW Queries to a file in cloud storage using the Data Transfer operator in the Modeler.

Attributes

Attribute	Description
Host	Host name of the ABAP Web server, without the protocols HTTP or HTTPS. <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px;"><p>→ Tip</p><p>To identify the host name, login to the SAP BW system with SAP GUI and run the transaction INA Testmonitor using the code <code>/nrstbit</code> in the upper left input field. To view the corresponding host and port to use, switch the protocol.</p></div>
Port	Port number of the ABAP web server.
Client	Client for the SAP BW system to use for login. If you don't provide a value, the application uses the default client of the SAP BW system.
Protocol	Protocol to use for connection: HTTP or HTTPS (default).
User	User name for the SAP BW (ABAP) system.
Password	Specified user's password for the SAP BW (ABAP) system.

Attribute	Description
HANA DB Connection ID	<p>If the underlying database is an SAP BW, SAP HANA, or SAP BW/4HANA system, the ID of the HANA_DB connection to the underlying database.</p> <p>The application requires this ID to optimize the Data Transfer operator by routing it through SAP HANA. To achieve this optimization, create another HANA_DB connection in SAP Data Intelligence Connection Management and use the connection ID as a value for this attribute.</p>
ABAP Connection ID	<p>ID of an ABAP connection that points to the same ABAP system.</p> <p>The application uses this value to enable delta load in the BW transfer so the data flow is through the ABAP subengine ODP operator. The connection must support Web RFC and use HTTP protocol.</p>
Language	<p>Two-digit ISO code for language. If you don't provide a value, the application uses the default login language of the BW (ABAP) system.</p>

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.8 CLOUD_DATA_INTEGRATION

The CLOUD_DATA_INTEGRATION connection type connects to and accesses information from systems that provide OData-based APIs for data integration with SAP Data Intelligence and SAP BW/4 HANA.

The following are examples of the SAP Cloud applications that provide OData-based APIs.

- SAP Fieldglass
- SAP C/4HANA Sales Cloud
- SAP C/4HANA Service Cloud
- SAP S/4HANA Cloud

Operations

The CLOUD_DATA_INTEGRATION connection allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.

- Preview content of remote objects in Metadata Explorer.
- Run rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Read data using SAP Data Intelligence Modeler.

Attributes

Attribute	Description
Host	Host name for accessing the cloud OData service.
Port	Port number for accessing the cloud OData service.
Protocol	Protocol to use for accessing the OData service: HTTP or HTTPS (default).
Service Path	Relative path to the Cloud Data Integration service endpoint, without host and port. Start the path with a forward slash character.
Root Path	Path to a Namespace or Provider to browse as root.
Authentication	Authentication method that the application must use. Options include the following: <ul style="list-style-type: none"> • Basic: Requires user name and password. • OAuth 2: Uses OAuth 2 credentials. • NoAuth: No authentication. • ClientCertificate: Uses client certificate for credentials.
User	User name for authentication. Applicable for basic authentication.
Password	Password for authentication for the specified user name. Applicable for basic authentication.
OAuth 2 Grant Type	Type of grant. Options include the following: <ul style="list-style-type: none"> • Client credentials • Password • Password with confidential client Applicable for OAuth2 authentication.
OAuth 2 Token Endpoint	Applicable for OAuth2 authentication.
OAuth 2 User	Applicable when you set the following attribute values: <ul style="list-style-type: none"> • GrantType = Password or Password with confidential client • Authentication = OAuth2
OAuth 2 Client ID	Applicable for OAuth2 authentication.
OAuth 2 Client Secret	Applicable for OAuth2 authentication.

Attribute	Description
OAuth 2 Scope	Applicable for OAuth2 authentication.
OAuth 2 Token Request Content Type	Value for the content-type HTTP header that the application must use when requesting a token. Applicable for OAuth2 authentication.
Require CSRF Header	Specify whether to require a CSRF (Cross-Site Request Forgery) header. The default value is true (on).

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.9 CPEM

SAP CPEM connection type connects to SAP Event Mesh, a capability of SAP Integration Suite. The connection uses the AMQP over a secure WebSocket (WSS) connection with OAuth 2.0 authentication.

Note

SAP Event Mesh was formerly known as SAP Cloud Platform Enterprise Messaging.

Operations

The CPEM connection type allows for the following operations:

- Receive messages published with SAP Event Mesh.
- Publish messages with SAP Event Mesh.

Attributes

Attribute	Description
Host	Host name or IP address of the SAP Event Mesh gateway.
OAuth 2 Token Endpoint	Token endpoint to use for OAuth 2 authentication.
OAuth 2 Client ID	Client's ID for OAuth 2 authentication.
OAuth 2 Client Secret	Client's secret for OAuth 2 authentication.
Use Client Certificate	Specify whether to use the X.509 client certificate for user authentication and authorization. The default is false (off).

Attribute	Description
Client Certificate	X.509 client certificate for user authentication and authorization. Required if you set Use Client Certificate to true (on).
Client Private Key	X.509 client private key when client certificate has a separate private key. Required if you set Use Client Certificate to true (on).

4.4.10 CPI

The CPI connection type provides an HTTPS connection to an SAP Integration service that runs on SAP Integration Suite.

Note

SAP Integration Suite was known as SAP Cloud Platform Integration Suite. SAP Integration Suite is a suite of services that run on the SAP Business Technology Platform.

Note

Client certificate and key are now supported.

Operations

The CPI connection type allows for the following operations:

- Trigger the graph run of iFlows (integration flows) in the connected CPI service.
- Provide the response from the graph run to other downstream operators in the SAP Data Intelligence Modeler.

Create the CPI connection type and use it in the CPI-PI iFlow operator in the Modeler. Obtain all information necessary to create a CPI connection from the administrator of the SAP Integration service.

Note

For the CPI connection type, SAP supports HTTP Basic Authentication with TLS only. HTTPS is HTTP with TLS.

Attributes

Attribute	Description
Host	The host name of the SAP Integration service, without protocol and port. Use the domain name or IP address as the host name.
Port	The port number of the SAP Integration service. If you leave this attribute blank, the application uses the HTTPS default port of 443.
User	The user name required by the application to authenticate at the target SAP Integration service.
Password	The password of the specified user to authenticate at the target SAP Integration service.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.11 DATASERVICES

The Dataservices connection type connects to a SOAP server and accesses information from an SAP Data Services Administration server.

Operations

The Dataservices connection type allows for the following operations:

- Browse SAP Data Services jobs.
- Run SAP Data Services jobs in the SAP Data Services Job Server.
- Connect to On-Premise Data Services through the SAP Cloud Connector.

The Dataservices connection type supports only HTTP connection protocol and the secEnterprise authentication type.

Note

The connection type allows you to connect to SAP Data Services on-premise systems of version 4.2 SP2 Patch 2 and later. For connections to SAP Data Services system versions earlier than version 4.2 SP2 Patch 2, you must enter the runtime parameters manually in the SAP Data Intelligence Modeler.

Attributes

Attribute	Description
Host	Host name of the SAP Data Services Web Service. Note In most cases, the host name is the SAP Data Services Management Console host. The host is usually the virtual host that has been configured in the Cloud Connector.
Port	Port number of the SAP Data Services Web Service. Note The port is usually the virtual port that has been configured in the Cloud Connector.
CMS Host	Host name of the SAP Data Services Central Management System (CMS).
CMS Port	Port number of the SAP Data Services CMS.
Connection protocol	Protocol to use for connection: HTTP or HTTPS (default).
Username	User name with which to authenticate in the SAP Data Services system.
Password	Password for the specified user name.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.12 DB2

The DB2 connection type connects to and accesses information from IBM DB2 databases. The application supports DB2 Universal Database (UDB) version 10.x and later.


Operations

The DB2 connection type allows for the following operation:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.

- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Version	Version of the IBM DB2 to which you're connecting: DB2 UDB 11.x (default) or DB2 UDB 10.x.
Host	Host name of the IBM DB2 server.
Port	Port number of the IBM DB2 server. The default is 50000.
Database name	Name of the IBM DB2 database to which you're connecting.
Use TLS	Specify whether to use TLS for an encrypted connection.
Server Certificate	DB2 Server certificate. To select, use the  icon. Applicable when you set Use TLS to true (on).
User	User name of the DB2 user who has privileges to connect to the database.
Password	Password for the specified DB2 user.
Additional session parameters	Enter additional session-specific parameters, if necessary.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.13 GCP_BIGQUERY

The GCP_BIGQUERY connection type connects to a Google Cloud platform BigQuery account.

Operations

The GCP_BIGQUERY connection type allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Publish remote objects in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Load data to a Google BigQuery table in Modeler.

ⓘ Note

the access to Google BigQuery sources without providing an ODBC driver is deprecated. An ODBC driver is required to access the features in Google BigQuery and for performance. For example, an ODBC driver is required for previewing content in Metadata Explorer and reading tables and views in the Modeler.

ⓘ Note

to use the GCP_BIGQUERY connection type, you must have the permission `bigquery.datasets.get` granted at the project level for your Google Cloud BigQuery account.

Prerequisites

Install an ODBC driver before you configure and use a GCP_BIGQUERY connection type.

Download the Magnitude Simba drivers for BigQuery from Google Cloud platform.

ⓘ Note

the steps in this topic refer to the driver version 3.0.0.1001. However, the steps remain the same for newer minor driver versions.

1. Select the Linux 32-bit and 64-bit (tar.gz) version:
`SimbaODBCDriverforGoogleBigQuery_3.0.0.1001-Linux.tar.gz.`
2. Create the vsolution area: `mkdir -p gcp_bigquery_vsolution/content/files/flowagent.`
3. Create vsolution manifest as `gcp_bigquery_vsolution/manifest.json`:

```
{
```

```

    "name": "vsolution_gcp_bigquery",
    "version": "1.0.0",
    "format": "2",
    "dependencies": []
  }

```

Note

to upload a new driver later, modify the version to upload a new vsolution. For example, 1.0.1, 2.0.0, and so on. Then modify your layering strategy appropriately.

4. Extract the downloaded compressed TAR archive:

```

tar -xzf SimbaODBCDriverforGoogleBigQuery_3.0.0.1001-Linux.tar.gz
tar -xzf SimbaODBCDriverforGoogleBigQuery_3.0.0.1001-Linux/
SimbaODBCDriverforGoogleBigQuery64_3.0.0.1001.tar.gz -C
gcp_bigquery_vsolution/content/files/flowagent/
cp SimbaODBCDriverforGoogleBigQuery_3.0.0.1001-Linux/
GoogleBigQueryODBC.did gcp_bigquery_vsolution/content/files/flowagent/
SimbaODBCDriverforGoogleBigQuery64_3.0.0.1001/lib/

```

5. configure the driver to display proper error messages as follows:

```

mv gcp_bigquery_vsolution/content/files/
flowagent/SimbaODBCDriverforGoogleBigQuery64_3.0.0.1001/ErrorMessage/en-
US gcp_bigquery_vsolution/content/files/flowagent/
SimbaODBCDriverforGoogleBigQuery64_3.0.0.1001/lib/en-US

```

Ensure that the `gcp_bigquery_vsolution/content/files/flowagent/`
`SimbaODBCDriverforGoogleBigQuery64_3.0.0.1001/lib` folder has the following structure:

```

lib/:
cacerts.pem en-US EULA.txt GoogleBigQueryODBC.did
libgooglebigqueryodbc_sb64.so
lib/en-US:
DSMessages.xml DSCURLHTTPClientMessages.xml ODBCMessages.xml
SimbaBigQueryODBCMessages.xml SQLEngineMessages.xml

```

6. Create a properties file named `gcp_bigquery.properties` at `gcp_bigquery_vsolution/content/files/flowagent/`. Ensure that there's a driver manager relative to the location of the properties file as follows:

```

GOOGLEBIGQUERY_DRIVERMANAGER=./
SimbaODBCDriverforGoogleBigQuery64_3.0.0.1001/lib/
libgooglebigqueryodbc_sb64.so

```

7. Compress the vsolution in the `gcp_bigquery_vsolution` directory as follows:

```

cd gcp_bigquery_vsolution
zip -r gcp_bigquery_vsolution.zip ./

```

Import the vsolution in System Management

Perform the following steps only after you install the ODBC driver:

Note

this process requires the Tenant Administrator role.

1. Open the SAP Data Intelligence System Management application.

2. Open the *Tenant* tab.
3. Select **+** , and then choose the new `gcp_bigquery_vsolution.zip` file.
4. After the import is complete, open the *Strategy* subtab, and then select *Edit*.
5. Select **+** , and then choose the imported solution `vsolution_gcp_bigquery-1.0.0`.
6. Select *Save*.

Note

for changes to take effect, restart the Flowagent application.

Limitations


When you use the GCP_BIGQUERY connection type, keep in mind the following limitations:

- When a table contains unsupported features, such as nested and repeated columns, or unsupported data types, SAP Data Intelligence can't consume data.
As a workaround, create a view without unsupported data type columns.
- When fetching metadata and performing data preview of partition tables, SAP Data Intelligence doesn't support a value of **True** for *Require Partition Filter*.
Ensure that you set the option *Require Partition Filter* to **False**.

Attributes

Attribute	Description
Project ID	Google BigQuery project ID.
Keyfile	Keyfile to upload that contains the access credentials for a service account.
Additional Regions	<p>List of additional regions for browsing datasets. Region examples include EU, us-west1, and asia-east1.</p> <p>By default, the application browses for datasets only from Google's default location. If you provide an invalid region, the check status still passes, but failures happen when you're browsing objects.</p>

Note

for more information about supported remote systems and data sources, see SAP Note [2693555](#) .

4.4.14 GCP_DATAPROC

The GCP_DATAPROC connection type connects to a Google Cloud Dataproc cluster. Google Cloud Dataproc is a scalable, metered, and managed Spark and Hadoop service. Clusters can scale up and down, as required.

Operations

The GCP_DATAPROC connection type allows for the following operations in the SAP Data Intelligence Modeler application:


- Submit Spar, PySpark, Hive, or SparkSQL jobs to a Dataproc cluster in a graph.
- Use the connection with the Submit Hadoop Job operator.

Attributes

To find all attribute values, look in the *Configuration* tab in the GCP console.

Attribute	Description
Project ID	Unique identifier for the Google Cloud Platform (GCP) project. The Project ID often consists of the project name and a random number. GCP projects are high-level groupings that manage APIs, billing, permissions, and more.
Cluster Name	Name of the specific Dataproc cluster for this connection.
Region	GCP-specific region in which the cluster is located, such as <code>europa-west3</code> .
Zone	Zone in which the cluster is located using the standard form, such as <code>europa-west3-b</code> .
Keyfile	Private key associated with the user or service account with Dataproc permissions.

📌 Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#) .

4.4.15 GCP_PUBSUB

The GCP_PUBSUB connection type connects to and accesses information from a Google Cloud Platform (GCP) publish/subscribe service.

Operations

The GCP_PUBSUB connection type allows for the following operations in SAP Data Intelligence Modeler:

- Create a topic on the Google Cloud Pub/Sub service and send messages to the topic using the GCP Pub/Sub Producer operator.
- Subscribe to a topic on the Google Cloud Pub/Sub service and receive messages from the topic using the GCP Pub/Sub Consumer operator.

Attributes

Attribute	Description
Project ID	Project ID on GCP to which the service account belongs.
Key file	Keyfile that contains the access credentials for a service account on the GCP.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.16 GCS

The GCS connection type connects to and accesses information from objects in Google Cloud Storage.

Operations

The GCS connection type allows for the following operations in Metadata Explorer:

- Browse remote objects.
- Obtain fact sheets of remote objects.
- Preview content of remote objects.
- Profile remote objects.
- Prepare data using the Preparation application.

- Save data from the Preparation application.
- Run rules.
- Extract metadata to the catalog.

In addition, the GCS connection type allows for the following operations in SAP Data Intelligence Modeler application:

- Read and write files using the Read File and Write File operators.
- Copy, rename, and remove files using the Copy File, Move File, and Remove File operators.

Attributes

Attribute	Description
Project ID	ID of the Google Cloud Storage project to which you want to connect.
Key file	Contents of the key file used for authentication.
Root Path	Optional root path name for browsing objects. Prefix the value with the forward slash character. For example, /My Folder/MySubfolder. If you specify a Root Path, the application prefixes the root path to any path used with this connection.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.17 HANA_DB

The HANA_DB connection type connects to and accesses information from tables and views in an SAP HANA database.

Operations

The HANA_DB connection type allows the following operations:

- Browse folders and files in the SAP HANA database.
- Obtain file metadata.
- Profile data.
- Preview data.
- Perform data preparation tasks.



Note

SAP recommends that you have DATA_ADMIN or CATALOG_READ system privileges or the SELECT_METADATA privilege on all schemas connected to HANA_DB.


Attributes

Attribute	Description
Host	Host name or IP address of the SAP HANA server. For SAP HANA Cloud, use the endpoint <i>without</i> the following element: :<port>.
Port	SQL port of the SAP HANA database: <ul style="list-style-type: none">• Single database: use 3<instance number>15.• Tenant database: check the port details by running the following SQL statement in your System DB: <pre>SELECT DATABASE_NAME, SERVICE_NAME, PORT, SQL_PORT FROM SYS_DATABASES.M_SERVICES</pre>• SAP HANA Cloud: use port 443 (provided at the end of the endpoint information of the SAP HANA Cloud instance).
Additional Hosts	List of host-port pairs to use as fallback hosts. The additional hosts are alternatives to the required Host and Port attributes.
Authentication	Authentication type to use to connect to the SAP HANA database. Authentication options include the following: <ul style="list-style-type: none">• Basic: requires a user name and password.• ClientCertificate: requires X.509 certification. When you select ClientCertificate, you also must complete the Client Certificate and the Client Private Key options.
User	User name for server authentication (Basic authentication).
Password	Password for the specified user (Basic authentication).

Attribute	Description
Blocked schemas	<p>List of database schemas to disable or hide from other applications.</p> <p>To enter a list, select + (<i>Add item</i>), and enter the names of the schemas to hide from other applications, such as the browse and publication tasks in Metadata Explorer.</p> <p>Other calls don't use the blocked schemas list, and the application shows the objects, including during lineage extraction in Metadata Explorer.</p> <p>Use the wildcard <code>*</code> to hide objects and datasets.</p> <div data-bbox="826 741 1385 1128" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>❖ Example</p> <ul style="list-style-type: none"> • <code>"_SYS*"</code> hides the top-level objects that start with <code>"_SYS"</code>. • <code>"CUSTOMER/*"</code> hides all of the objects in the CUSTOMER object. • <code>"CUSTOMER/PROSPECTS*"</code> hides all objects in CUSTOMER objects that begin with <code>"PROSPECTS"</code>. • <code>"CUSTOMER_EUROPE"</code> hides the top-level objects that are an exact match to <code>"CUSTOMER_EUROPE"</code>. </div>
Use TLS	<p>Specify whether to use TLS to connect to the server.</p> <p>For SAP HANA Cloud, TLS must be enabled (toggled on).</p> <p>TLS is enabled by default for all SAP HANA connections.</p>
Validate host certificate	<p>Specify whether to validate the server certificate.</p> <div data-bbox="826 1413 1385 1554" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>📘 Note</p> <p>Validate host certificate is applicable only when you set Use TLS to true (on).</p> </div>
Hostname in certificate	<p>Specify the host name to validate in the certificate.</p> <div data-bbox="826 1650 1385 1778" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>📘 Note</p> <p>Hostname in certificate is applicable only when you set Validate Host Certificate to true (on).</p> </div>

Attribute	Description
Client certificate	<p>File that contains the X.509 client certificate in your repository. The file must contain only the client certificate and not be bundled with the client private key.</p> <p>To browse for the file, select  (Browse).</p> <p>Required only when you choose ClientCertificate for Authentication.</p>
Client Private Key	<p>File that contains the X.509 client private key in your repository. The file must contain only the client private key and not be bundled with the client certificate.</p> <p>To browse for the file, select  (Browse).</p> <p>Required only when you choose ClientCertificate for Authentication.</p>

Note

For information about supported versions, see SAP Note [2693555](#) .

4.4.18 HANA_XS

The HANA_XS connection type connects to and accesses information from tables and views in SAP HANA databases.

Operations

The HANA_XS connection type allows for the following operations:

- Browse SAP HANA flowgraphs.
- Run SAP HANA flowgraphs.

Attributes

Attribute	Description
Host	Host name or IP address of the SAP HANA server.
Port	Port number of the SAP HANA server.

Attribute	Description
Protocol	Protocol to use for connection. Options are HTTP or HTTPS. The default value is HTTPS.
User	User name to authenticate the server.
Password	Password for the specified user.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.19 HDFS

The HDFS connection type connects to and accesses information from objects in an HDFS (Hadoop Distributed File System) server.

Operations

The HDFS connection type allows for the following operations:

- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Save data from the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Read and write files using the Read File and Write File operators.
- Rename and remove files using the Move File and Remove File operators.

The HDFS connection type also supports the following features:

- Use RPCs (remote procedure calls) to extend an HDFS connection to use WebHDFS and SWebHDFS.
- Choose from the following authentication types: Kerberos, basic, or simple.

Attributes

Attribute	Description
Host	Host name or IP address of the HDFS namenode.
Port	Port number of the HDFS namenode. If you leave this attribute blank, the application provides a default value based on the selected protocol.
Additional Hosts	List of secondary hostnames or IP addresses, which is required for HDFS High Availability (HA) configurations.
Protocol	List of secondary hostnames or IP addresses, which is required for HDFS High Availability (HA) configurations.
	<div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px;"> <p>Note</p> <p>If you use HDFS in a network other than the SAP Data Intelligence cluster, you can use only webhdfs or swebhdfs connections, but not RPC. Instead, run HTTPFS proxy on HDFS and use only HTTPFS port to reach the HDFS server.</p> </div>
Authentication Type	<p>Authentication mechanism that the application uses to connect to the specified HDFS. Options include the following mechanisms:</p> <ul style="list-style-type: none"> • simple • basic • kerberos <p>If you choose <i>kerberos</i>, also complete the following parameters.</p> <ul style="list-style-type: none"> • krb5.conf File • Keytab File
User	User name to access the specified HDFS. If you select kerberos for authentication, provide the user principal.
Root Path	<p>Optional root path name for browsing objects. Start the value with a forward slash. For example, <code>/My Folder/MySubfolder</code>.</p> <p>If you've specified the Root Path, then any path used with this connection is prefixed with the Root Path.</p>
krb5.conf File	krb5.conf file that contains at a minimum, the realm and corresponding kdc for the environment. Applicable when Authentication Type = kerberos.
Keytab File	Keytab file for the principal named in the User attribute. Applicable when Authentication Type = kerberos.
Impersonation	Specify whether to impersonate the connected users (act on-behalf of). The default setting is false (off).

Attribute	Description
Custom Parameters	List of HDFS custom parameters.
Proxy	Optional. Connection-specific proxy server. Specify type, host, and port. The available types are HTTP (default) and SOCKS.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.20 HDL_DB

The HDL_DB connection type connects to and accesses information from SAP HANA Cloud data lake databases.

Operations

The SAP HANA Cloud data lake database connection allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- View fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Save data from the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Authentication	Authentication method to use for connection to SAP HANA Cloud data lake databases: <ul style="list-style-type: none">• Basic: Requires username and password.• ClientCertificate: Requires username and X509 client certification.
Host	Host name of the SAP HANA Cloud data lake database server.
Port	Port number of the SAP HANA Cloud data lake database server.
User	Username of the SAP HANA Cloud data lake user. The user must have privileges to connect to the database name.
Password	Password for the specified SAP HANA Cloud data lake database user.
Client Certificate	File that contains the X.509 client certificate in your repository. File must contain only the client certificate and can't be bundled with the client private key. To browse for the file, choose Browse . Required only when you choose ClientCertificate for Authentication .
Client Private Key	File that contains the X.509 client private key in your repository. File must contain only the client private key and can't be bundled with the client certificate. To browse for the file, choose Browse . Required only when you choose ClientCertificate for Authentication .
Client Private Key Password	Password for the X.509 Client Private Key.
Additional session parameters	Session-specific variables, if necessary.

Datatype Conversion

The application converts data types from SAP HANA data lake database to an agnostic set of data types. The application supports only SAP HANA data lake database datatypes that have a corresponding datatype. The following table shows the data type conversions.

SAP HANA data lake Database Data Type	Mapped Data Type
BIGINT	DECIMAL(20, 0)
UNSIGNED BIGINT	DECIMAL(20, 0)

SAP HANA data lake Database Data Type	Mapped Data Type
BIT	INTEGER(4)
DECIMAL(p,s)	DECIMAL(p, s)
DOUBLE	FLOATING(8)
FLOAT	FLOATING(4)
INT	INTEGER(4)
UNSIGNED INT	INTEGER(4)
INTEGER	INTEGER(4)
UNSIGNED INTEGER	INTEGER(4)
MONEY	DECIMAL(19,4)
NUMERIC(p,s)	DECIMAL(p, s)
REAL	FLOATING(4)
SMALLINT	INTEGER(4)
SMALLMONEY	DECIMAL(10,4)
TINYINT	INTEGER(4)
UNIQUEIDENTIFIER	N/A
VARCHAR(s)	STRING(s)
SYSNAME	STRING(30)
DATE	DATE
DATETIME	DATETIME
SMALLDATETIME	DATETIME
TIME	TIME
TIMESTAMP	DATETIME
BINARY	N/A
VARBINARY	N/A
BLOB	LARGE_BINARY_OBJECT
CLOB	LARGE_CHARACTER_OBJECT
IMAGE	LARGE_BINARY_OBJECT
LONG BINARY	LARGE_BINARY_OBJECT
TEXT	LARGE_CHARACTER_OBJECT

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.21 HDL_FILES

The HDL_Files connection type connects to and accesses information from file storage on SAP HANA Data Lake.

Operations

The SAP HANA Data Lake database connection allows for the following operations in Metadata Explorer:

- Browse remote objects.
- View fact sheets of remote objects.
- Preview content of remote objects.
- Profile remote objects.
- Prepare data using the Preparation application.
- Save data from the Preparation application.
- Run rules.
- Extract metadata to catalog.

Attributes

Attribute	Description
Keystore File	Client keystore to use. Either select the file in P12 format-binary or enter the Base64-encoded value of P12 file content.
Endpoint	SAP HANA Data Lake File endpoint.
Root Path	Optional root path name for browsing. Start the path with a forward slash character, for example, /MyFolder/MySubfolder. The application prefixes any path used with this HDL_Files connection with the specified root path.
Keystore Pwd	Client keystore Password. The connection fails without a password.
Impersonate	Specify whether to impersonate the current user (act on-behalf of). The default setting is false (off).

Example

The following sample code generates a keystore from a certificate and key to access HDL files:

Sample Code

```
openssl pkcs12 -export -inkey client.key -in client.crt -out keystore.p12
```

4.4.22 HTTP

The HTTP connection type connects to and accesses information from servers over HTTP or HTTPS.


Operations

The HTTP connection type allows for accessing an arbitrary HTTP endpoint.

Attributes

Attribute	Description
Host	Host name or IP address of the HTTP server.
Port	Port number for the HTTP server.
Protocol	Protocol to use for connecting to the server: HTTP or HTTPS. The default value is HTTPS.
Path	Path prefix to the API endpoint. For example, <code>/api/xyz</code> .
Authentication	<p>Authentication method to use when connecting to the server. Options include the following:</p> <ul style="list-style-type: none">• NoAuth• Basic• OAuth2 <p>If you select Basic, then provide values to the following attributes:</p> <ul style="list-style-type: none">• User• Password <p>If you select OAuth2, then provide values to the following attributes:</p> <ul style="list-style-type: none">• OAuth 2 Grant Type• OAuth 2 Token Endpoint• OAuth 2 Client ID• OAuth 2 Client Secret• OAuth 2 Scope• OAuth 2 Resource• OAuth 2 Token Request Content Type

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#) .

4.4.23 IMAP

The IMAP connection type connects to and accesses information from an IMAP (Internet Message Access Protocol) server to receive emails.

The IMAP connection follows the protocol standard defined in RFC 350.

Operations

The IMAP connection type allows you to use the Receive Email operator in the SAP Data Intelligence Modeler application.

Attributes

Attribute	Description
Host	Host name of the IMAP server.
Port	Port number of the IMAP server.
Authorization Method	Authorization method to use. Options include the following: <ul style="list-style-type: none">• Password: requires username and password for authorization.• ms_client_credentials: requires Microsoft Exchange client credentials.
Use Client Certificate	Specify whether to use Microsoft Exchange client certificate instead of the client secret. Applicable when Authorization Method is ms_client_credentials.
User	Name of the user who accesses the IMAP server.
Password	Password of the specified user. Applicable when Authorization Method is password.
Authority	Authority for Microsoft Exchange client certificate authentication.
Scope	Scope for Microsoft Exchange client certificate authentication. Applicable when Authorization Method is ms_client_credentials.
ClientId	Client ID for Microsoft Exchange authentication. Applicable when Authorization Method is ms_client_credentials.
ClientSecret	Client secret for Microsoft Exchange authentication. Applicable when Authorization Method is ms_client_credentials.

Attribute	Description
ClientCertificate	Client certificate for Microsoft Exchange client certificate authentication. Applicable when Use Client Certificate is true (on).
ClientPrivateKey	Client private key for Microsoft Exchange client certificate authentication. Applicable when Use Client Certificate is true (on).
Use TLS	Specify whether to use TLS for connecting to the IMAP server. The default is false (off).

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.24 INFORMATION_STEWARD

The Information_Steward connection type connects to and accesses information from an SAP Information Steward Administration server.

Note

The connection is shown only in the Metadata Explorer when importing rules; it's not shown when browsing connections.

Operations

The Information Steward connection type allows the following operations:

- View connection status in Connection Management.
- Connect using HTTP and HTTPS protocol.
- Authenticate using secEnterprise type.
- View SAP Information Steward projects in Metadata Explorer.
- Import SAP Information Steward rules and bindings in Metadata Explorer.
- Connect to on-premise SAP Information Steward via SAP Cloud Connector.

Attributes

Attribute	Description
Host	Host name of the SAP Information Steward Web Service host.
Port	Port number of the Information Steward Web Service.
CMS Host	Host name of the SAP Central Management System (CMS).
CMS Port	Port number of the SAP Central Management System.
Connection protocol	Protocol to use for connection. Options are as follows: HTTP or HTTPS. The default value is HTTPS.
Username	User name to authenticate in the SAP Information Steward system.
Password	Password of the specified user to authenticate in the SAP Information Steward system.

There are additional options when you connect on cloud instances. For more information, see [Using SAP Cloud Connector Gateway \[page 223\]](#).

4.4.25 KAFKA

The KAFKA connection type connects to and accesses information from an Apache Kafka cluster.

Operations

The KAFKA connection allows for the following operations:

- Consume messages from a list of Kafka topics using the Kafka Consumer operator.
- Produce messages to a Kafka topic using Kafka Producer operator.

Note



For information about supported versions, see SAP Note [2693555](#).

Attributes

Attribute	Description
Kafka Brokers	Comma-separated list of Kafka brokers using the format <code>host:port</code> .

Attribute	Description
Group ID	Group ID contains all information for a consumer that is part of the group.
Use TLS	<p>Specify whether to connect to Kafka using TLS protocol. The default setting is true (on).</p> <p>When turned on, complete the following additional attributes:</p> <ul style="list-style-type: none"> • Insecure skip verify • Use Certificate Authority • Client Certificate or Client Certificate Path • Client Key or Client Key Path
Authentication	<p>Specify whether to authenticate using a username and password. The default setting is true (on).</p> <p>When turned on, also complete the following attributes:</p> <ul style="list-style-type: none"> • SASL Authentication • Kafka SASL Username • Kafka SASL Password
Insecure skip verify	<p>Specify whether the KAFKA pipeline operator checks the TLS server certificates for authentication. The default setting is false (off).</p> <div data-bbox="804 1151 1394 1305" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>⚠ Caution</p> <p>Setting this attribute to true (on) makes the connectivity to Kafka less secure.</p> </div>

Attribute	Description
Credential Storage Inline	<p>Location where you store your KAFKA credentials.</p> <p>False (off): you store KAFKA credentials in the SAP Data Intelligence repository. False (off) is the default setting. When turned off, set the following attributes:</p> <ul style="list-style-type: none"> • Client Certificate Path • Client Key Path • Certificate Authority Path • Kerberos Config Path <p>True (on): you store KAFKA credentials in the connection. When turned on, set the following attributes:</p> <ul style="list-style-type: none"> • Client Certificate • Client Key • Certificate Authority • Kerberos Config
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <p>Storing credentials in the repository exposes them to other SAP Data Intelligence applications, which is less secure than storing them in the secure connection.</p> </div>	
SASL Authentication	<p>Simple Authentication and Security Layer (SASL) authentication mechanism to use. Options include the following:</p> <ul style="list-style-type: none"> • PLAIN: plaintext password defined in RFC 4616. • SCRAM-256: Salted Challenge Response Authentication Mechanism (SCRAM) password-based challenge-response authentication from a user to a server. Based on SHA 256. • SCRAM-512: SCRAM password-based challenge-response authentication from a user to a server. Based on SHA 512. • GSSAPI: Generic Security Service Application Program Interface authentication applicable for Kerberos V5. <p>Applicable when you set Authentication to true (on).</p>
Kafka SASL Username	<p>Kafka SASL connection username.</p> <p>Applicable when you set Authentication to true (on).</p>
Kafka SASL Password	<p>Kafka SASL connection password.</p> <p>Applicable when you set Authentication to true (on).</p>

Attribute	Description
Use Certificate Authority	<p>Specify whether to use a TLS certificate authority (CA) file to connect to Kafka:</p> <ul style="list-style-type: none"> • True (on): connects to Kafka using TLS CA. • False (off): connects to Kafka without using TLS CA. <p>Applicable when you set Authentication to true (on).</p>
Certificate Authority	<p>Certificate authority file to upload to the secure connection.</p> <p>To select the file, use the  <i>icon</i>.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = True (on) • Use TLS = True (on) • Use Certificate Authority = True (on)
Certificate Authority Path	<p>Path to the SAP Data Intelligence repository that stores the applicable certificate authority file.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = False (off) • Use TLS = True (on) • Use certificate authority = True (on)
Client Certificate	<p>Client certificate file to upload to the secure connection. To select the file, use the  <i>icon</i>.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = True (on) • Use TLS = True (on)
Client Certificate Path	<p>Path to the SAP Data Intelligence repository that stores the applicable client certificate file.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = False (off) • Use TLS = True (on)

Attribute	Description
Client Key	<p>Client key certificate file to upload to the secure connection. To select the file, use the  <i>icon</i>.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = True (on) • Authentication = True (on) • SASL Authentication = GSSAPI <p>For Java keystore, use the <code>-keyalg RSA</code> constraint.</p> <div data-bbox="804 678 1394 831" style="background-color: #f0f0f0; padding: 5px;"> <p>Note</p> <p>The supported encryption key types are: rsa, ecdsa, and ed25519.</p> </div>
Client Key Path	<p>Path to the SAP Data Intelligence repository that stores the applicable client key file, which is loaded with the client certificate file.</p> <p>For Java keystore, use the <code>-keyalg RSA</code> constraint.</p> <div data-bbox="804 1037 1394 1189" style="background-color: #f0f0f0; padding: 5px;"> <p>Note</p> <p>The supported encryption key types are: rsa, ecdsa, and ed25519.</p> </div> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = False (off) • Use TLS = False (off)
Kafka Kerberos Service Name	<p>Service name defined for the Kafka server.</p> <p>Applicable when SASL Authentication = GSSAPI.</p>
Kafka Kerberos Realm	<p>Realm defined for the Kafka Kerberos server.</p> <p>Applicable when SASL Authentication = GSSAPI.</p>
Authentication Type For Kafka/GSSAPI	<p>Authentication type for Kafka/GSSAPI:</p> <ul style="list-style-type: none"> • User/password • Keytab file authentication <p>Applicable when SASL Authentication = GSSAPI.</p>
Kafka SASL Kerberos Password	<p>Password for Kafka SASL Kerberos connection.</p> <p>Applicable when Authentication type for Kafka/GSSAPI = User/password.</p>

Attribute	Description
Kerberos Config	<p>Kerberos configuration file (krb5.conf) to upload to the secure connection. Use the  icon.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = True (on) • Authentication = True (on) • SASL Authentication = GSSAPI • Authentication type for Kafka/GSSAPI = Keytab file authentication
Kerberos Config Path	<p>Path to the SAP Data Intelligence repository that stores the applicable Kerberos configuration file (krb5.conf).</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = False (off) • Authentication = True (on) • SASL Authentication = GSSAPI • Authentication type for Kafka/GSSAPI = Keytab file authentication
Keytab File	<p>Keytab file to upload to the secure connection. Use the  icon.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = True (on) • Authentication = True (on) • SASL Authentication = GSSAPI • Authentication type for Kafka/GSSAPI = Keytab file authentication
Keytab File Path	<p>Path to the keytab file for Kafka client.</p> <p>Applicable with the following attribute settings:</p> <ul style="list-style-type: none"> • Credential Storage Inline = False (off) • Authentication = True (on) • SASL Authentication = GSSAPI • Authentication type for Kafka/GSSAPI = Keytab file authentication

4.4.26 MSSQL

The MSSQL connection type connects to and accesses information from Microsoft SQL Server databases (MSSQL). The connection type supports MSSQL version 2012 and later versions.

Operations

The MSSQL connection type allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Version	Microsoft SQL Server version: 2012 and later versions..
Subtype	Microsoft SQL Server subtype. The application supports only SQL Server on Premise.
Host	MSSQL server host name.
Port	MSSQL server port number.
Use TLS	Specify whether to use TLS encryption to connect to MSSQL. The default setting is false (off).
Database name	Name of the MSSQL database with which to establish a connection.
User	User name for the MSSQL database. The user must have the privilege to connect to the database.
Password	Password for the specified MSSQL user.

Attribute	Description
Additional session parameters	Additional session-specific parameters, if necessary.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.27 MYSQL

The MySQL connection type connects to and accesses information from Oracle MySQL server. The application supports MySQL version 5.5 and later versions.

Operations

The MYSQL connection type allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Version	MySQL server version. The default value is MySQL 8.x.
Host	MySQL server host name.
Port	MySQL server port number. The default value is 3306.
Database name	Name of the MySQL database with which to connect.

Attribute	Description
User	User name of the MySQL database. The user must have privileges to connect to the database.
Password	Password of the specified MySQL user.
Additional session parameters	Additional session-specific parameters, as necessary.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.28 ODATA

The OData connection type connects to and accesses information from OData RESTful APIs. SAP Data Intelligence supports connection to OData Version V2 and V4.

Operations

The ODATA connection type allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Run rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Read OData resources in Modeler.

Attributes

Attribute	Description
URL	URL for the OData Service.
Version	OData Version: V2 or V4.

Attribute	Description
Authentication	Method for authentication: <ul style="list-style-type: none"> • Basic • OAuth2 • NoAuth • ClientCertificate
User	User name to authenticate. Applicable for Basic authentication.
Password	Specified user password. Applicable for Basic authentication.
Require CSRF Header	Specify whether to require a CSRF (cross-site request forgery) header.
Client Certificate	X.509 client certificate to use for ClientCertificate authentication.
Provide Client Private Key in separate file	Specify whether to use the X.509 client private key in a separate file.
Client Private Key	X.509 client private key when you set Provide Client Private Key in separate file to true (on).
Use Client Private Key Password	Specify whether to use an X.509 client private key that is protected by a password.
Client Private Key Password	Password for the X.509 client private key when you set Use Client Private Key Password to true (on).
OAuth2 GrantType	OAuth2 grant type. Applicable when you use OAuth2 authentication. The following list shows the supported grant types: <ul style="list-style-type: none"> • Password with confidential client • Client credentials • Password • Password with confidential client • SAML2 bearer
OAuth2 Token Endpoint	OAuth2 token endpoint. Applicable when you use OAuth2 authentication.
OAuth2 User	OAuth2 user name. Applicable when you use OAuth2 authentication.
OAuth2 Password	OAuth2 password for the specified OAuth2 user. Applicable when you use OAuth2 authentication.
OAuth2 ClientId	OAuth2 client ID. Applicable when you use OAuth2 authentication.
OAuth2 Client Secret	OAuth2 client secret. Applicable when you use OAuth2 authentication.
OAuth2 Scope	OAuth2 scope. Applicable when you use OAuth2 authentication.

Attribute	Description
OAuth2 Resource	OAuth2 resource. Applicable when you use OAuth2 authentication.
OAuth2 Response Type	OAuth2 response type. Applicable when you use OAuth2 authentication.
OAuth2 Token Request Content Type	<p>OAuth2 token request content type. Applicable when you use OAuth2 authentication. Options include the following:</p> <ul style="list-style-type: none"> • url_encoded: the OAuth2 token request parameters are url-encoded and included in the HTTP request body (default). • json: the OAuth2 token request parameters are in JSON format and included in the HTTP request body.
OAuth2 API Endpoint	OAuth2 token endpoint. Applicable when you use OAuth2 authentication and saml2_bearer grant type.
OAuth2 UserId	OAuth2 user ID. Applicable when you use OAuth2 authentication and saml2_bearer grant type.
OAuth2 CompanyId	OAuth2 company ID. Applicable when you use OAuth2 authentication and saml2_bearer grant type.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.29 OPEN_CONNECTORS

The Open_Connectors connection type connects to and accesses data from Open Connectors, a core capability of SAP Integration Suite.

Prerequisites

Before you create an Open_Connectors connection type, configure the Open Connectors service by performing the following steps:

1. Configure an Open Connectors service. For information about configuring an Open Connectors trial version, see [Enable SAP BTP Open Connectors in Trial](#).
2. Open the Connectors tab in the *SAP BTP Open Connectors* landing page.
3. Select **Authenticate > Create Instance**.
4. Note the following information about your Open Connectors instance so that you can complete the Open_Connectors configuration in SAP Data Intelligence Connection Management:
 - Users secret
 - Organization secret

- Instance ID

SAP Data Intelligence supports all open connectors from the following hub types:

- CRM
- DB
- General
- Marketing
- Social

For more information about the connectors from each of the hub types, see the [SAP Open Connectors](#) documentation.

⚠ Restriction

These hub types are experimental. 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.

Operations

The Open_Connectors connection allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Run rules in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.

Attributes

Attribute	Description
Open Connectors Instance ID	ID of an authenticated connector instance. Each authenticated connector instance is identifiable by its unique ID.

Attribute	Description
Open Connectors API Base URL	API Base URL that accesses Open Connectors. For more information about Open_ Connectors API Base URL, see the article Base URL .
User Secret	User Secret for Open Connectors.
Organization Secret	Organization Secret for Open Connectors.

ⓘ Note

If the article doesn't open initially, enter "Base URL" in the Search field of the page.

ⓘ Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.30 OPENAPI

The OpenAPI connection type connects to and accesses information from a server according to an OpenAPI 2.0 specification.

Attributes

Attribute	Description
Host	Service's host name.
Port	Service's port number.
Protocol	Protocol to use for connection: HTTP or HTTPS.
Basepath	Service's base path.

Attribute	Description
Authentication Type	<p>Options include the following:</p> <ul style="list-style-type: none"> • NoAuth • Basic • OAuth2 • ApiKey <p>For Basic, provide values for the following attributes:</p> <ul style="list-style-type: none"> • User: authentication user • Password: authentication password <p>For OAuth2, provide values for the following attributes:</p> <ul style="list-style-type: none"> • OAuth2 Grant Type • OAuth2 Token Endpoint • OAuth2 User • OAuth2 Password • OAuth2 Client Id • OAuth2 Client Secret • OAuth2 Authorization Server: Server for authorization code grant type • OAuth2 Scope • OAuth2 Resource <p>For ApiKey, provide values for the following attributes:</p> <ul style="list-style-type: none"> • ApiKey Name: name of the ApiKey parameter • ApiKey Type: type of the ApiKey (header or query) • ApiKey Value: value of the ApiKey, if ApiKey authentication is chosen.
TLS skips verify	Specify whether the client skips verifying the server certificate chain or the host name.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.31 ORACLE

The Oracle connection type connects to and accesses information from an Oracle database. The versions supported for connection are Oracle 18c and Oracle 19c.

Note

The support for version 10g, 11g, and 12c is deprecated.

Prerequisites

Before you configure an Oracle connection, install the Oracle Instant Client for Linux x86-64 at <https://www.oracle.com/database/technologies/instant-client/linux-x86-64-downloads.html>. Select the Basic Light Package from the 19.17.0.0 version. Then perform the following steps:

1. Create the vsolution area:
`mkdir -p oracle_vsolution/content/files/flowagent.`
2. Create a vsolution manifest as `oracle_vsolution/manifest.json` using the following sample code:

Sample Code

```
{
  "name": "vsolution_oracle",
  "version": "1.0.0",
  "format": "2",
  "dependencies": []
}
```

Note

To upload a new driver later, modify the version (for example, 1.0.1, 2.0.0, and so on) to upload a new vsolution. Then modify your layering strategy appropriately.

3. Extract the zip file and create the applicable symbolic links using the following command:

```
unzip instantclient-basiclite-linux.x64-19.17.0.0.dbru.zip -d
oracle_vsolution/content/files/flowagent/
cd oracle_vsolution/content/files/flowagent/instantclient_12_2
ln -s libclntsh.so.12.1 libclntsh.so.12
ln -s libclntsh.so.12 libclntsh.so
ln -s libclntshcore.so.12.1 libclntshcore.so.12
ln -s libclntshcore.so.12 libclntshcore.so
ln -s libocci.so.12.1 libocci.so.12
ln -s libocci.so.12 libocci.so
cd ../../../../..
```

Note

If your version isn't 19.17.0.0, refer to the readme file that comes with the installation package, and follow the steps for creating symbolic links. The zip file name depends on the driver version downloaded.

4. Create a properties file `oracle_vsolution/content/files/flowagent/oracle.properties` with Oracle instant client path relative to the location of the property file:

```
ORACLE_INSTANT_CLIENT=./instantclient_19_17
NLS_LANG=AMERICAN_AMERICA.UTF8
```

5. Configure TLS by performing the following substeps:
 1. Create `orapki` directory: `mkdir -p oracle_vsolution/content/files/flowagent/orapki.`
 2. Find and copy `oraclepki.jar`, `osdt_core.jar`, and `osdt_cert.jar` to the `orapki` directory.

Note

Find these `.jar` files in the official Oracle JDBC website (<https://www.oracle.com/database/technologies/appdev/jdbc-downloads.html>). Download the zipped JDBC driver (`ojdbc8.jar`)

and Companion Jars for version 19.17.0.0 (https://download.oracle.com/otn-pub/otn_software/jdbc/1917/ojdbc8-full.tar.gz). These .jar files are located at the root of the extracted ZIP file.

3. Add `ORACLE_ORAPKI_PATH=./orapki` to `oracle.properties`.
6. Compress the vsolution from the `oracle_vsolution` directory.

Sample Code

```
cd oracle_vsolution
zip -y -r oracle_vsolution.zip ./
```

7. Import the vsolution using System Management by performing the following substeps:

Note

You must have the Tenant Administrator role.

1. Start the System Management application from the Launchpad.
2. Open the *Tenant* tab.
3. Select **+** (*Add Solution*), and choose the `oracle_vsolution.zip` file.
4. Open the *Strategy* tab and select **✎** (*Edit*).
5. Select **+** (*Add Solutions*).
6. Choose the `vsolution_oracle-1.0.0` solution and select *Add*.
7. Select *Save*.

Note

For changes to take effect, restart the Flowagent application.

Operations

The Oracle connection allows the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Publish remote objects in Metadata Explorer.
- Profile remote objects in Metadata Explorer.
- Prepare data using the Preparation application in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.

- Run native SQL DDL or DML statements in Modeler.

Attributes

Attribute	Description
Version	Oracle database version to which you're connecting, such as Oracle 19c.
Host	Host name of the Oracle database server.
Port	Port number of the Oracle database server. The default port is 1521.
SID or Service Name	Database instance or alias to which you establish a connection.
Validate host certificate	Specify whether to validate the server certificate. The default is false (off).
Hostname in certificate	Host name to validate in certificate.
User	Oracle database user with privileges to connect to the database instance.
Password	Password for the specified Oracle database user.
Additional session parameters	Enter any session-specific variables, if necessary.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.32 OSS

The OSS connection type connects to and accesses information from an Alibaba Cloud Object Storage Service (OSS).

Operations

The OSS connection allows for the following operations in Metadata Explorer:

- Browse remote objects.
- Obtain fact sheets of remote objects.
- Preview content of remote objects.
- Profile remote objects.
- Prepare data using the Preparation application.

- Save data from the Preparation application.
- Run rules.
- Extract metadata to catalog.

The OSS connection allows for the following operations in the Modeler:

- Read and write files in Modeler with the Read File and Write File operators.
- Copy, rename, and remove files in Modeler with the Copy File, Move File, and Remove File operators.

Attributes

Attribute	Description
Endpoint	OSS server endpoint URL. The protocol prefix isn't required. For example, "oss.aliyuncs.com".
Protocol	<p>Protocol to use. Options include the following:</p> <ul style="list-style-type: none"> • HTTPS (default) • HTTP <p>The Protocol value overwrites the value from the Endpoint attribute, if you've already set it.</p>
Region	Region over which the connection authenticates and operates. For example, "oss-cn-hangzhou".
Access Key	User's Access Key ID for authentication.
Secret Key	User's Secret Access Key for authentication.
Root Path	<p>Optional root path name for browsing. The value must start with a forward slash character. For example, /My Folder / MySubfolder.</p> <p>If you specify a Root Path, then the application prefixes any path used with this connection with the Root Path.</p>

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.33 POSTGRESQL

The PostgreSQL connection type connects to and accesses information from PostgreSQL databases. PostgreSQL versions 13.X and 14.X are supported.

Operations

The PostgreSQL connection type allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Publish remote objects in Metadata Explorer.
- Read tables and views in the Modeler.
- Read data from SQL queries in the Modeler.
- Run native SQL DDL/DML statements in the Modeler.

The PostgreSQL operator reads a table from any database that provides an Open Database Connectivity (ODBC) connector. It uses the Flowagent subengine to connect to the database. To establish an ODBC connection, provide a Data Source Name (DSN) that specifies the DSN properties. You must define the DSN in a file using the UNIX standard.

Attributes

Attribute	Description
Version	The PostgreSQL version. Supported PostgreSQL versions are 13.X or 14.X. Version 14.X is the default value.
Host	The host name of the PostgreSQL server.
Port	The port number of the PostgreSQL server. The default port number is 5432.
Database name	The PostgreSQL database with which to connect.
Use TLS	Specifies whether to enable a TLS encrypted connection. The default setting is false (no).
Validate host certificate	Specifies whether to validate the server certificate. The default setting is false (no).
Validate hostname	Specifies whether to validate the hostname in the server certificate. The default setting is false (no).
Authentication	The authentication method to use for connection to the PostgreSQL database: <ul style="list-style-type: none">• Basic: Requires the PostgreSQL user name and password.• ClientCertificate: Requires PostgreSQL user name and X.509 client certificate.

Attribute	Description
User	The PostgreSQL user name who has privileges to connect to the database instance.
Password	The password of the specified PostgreSQL user. Applicable when you set Authentication to Basic.
Client Certificate	The X.509 client certificate. Applicable when you set Authentication to ClientCertificate.
Client Private Key	The X.509 client private key. Applicable when you set Authentication to ClientCertificate.
Client Private Key Password	The password for the X.509 client private key. Applicable when you set Authentication to ClientCertificate.

Note

The application converts NaN (Not a Number) values to 0 when present in decimal or numeric columns. For running graphs, the application writes OID (Object Identifier) data type values as -1 when out of the range of [-2147483648 through +2147483647]. The preview displays the original data.

4.4.34 REDSHIFT

The Redshift connection type connects to and accesses information from Amazon Redshift databases.

Prerequisites

Before you configure a Redshift connection, install the Amazon Redshift ODBC driver (64-bit .rpm version) for *Linux* operating systems:

1. Create the vsolution area:

```
mkdir -p redshift_vsolution/content/files/flowagent/redshift
```
2. Create a vsolution manifest as `redshift_vsolution/manifest.json`:

Sample Code

```
{
  "name": "vsolution_redshift",
  "version": "1.0.0",
  "format": "2",
  "dependencies": []
}
```

Note

To upload a new driver later, modify the version, such as 1.0.1, 2.0.0, and so on, to upload a new vsolution. Then modify your layering strategy appropriately.

3. Extract the downloaded RPM file based on your operating system:
 - On Windows, extract the RPM file using a file archiver.
 - On Linux, install the package or extract it using 'rpm2cpio' tool:

Sample Code

```
'rpm2cpio AmazonRedshiftODBC-64-bit-<version>.x86_64.rpm | cpio -idmv'
```

4. Copy the extracted files to vsolution using the following code:

Sample Code

```
cp <extracted_location>/opt/amazon/redshiftdbc/lib/64/  
libamazonredshiftdbc64.so redshift_vsolution/content/files/flowagent/  
redshift/  
cp <extracted_location>/opt/amazon/redshiftdbc/lib/64/  
AmazonRedshiftODBC.did redshift_vsolution/content/files/flowagent/redshift/
```

5. Configure the driver to display proper error messages using the following code:

Sample Code

```
cp -R <extracted_location>/opt/amazon/redshiftdbc/ErrorMessage/en-US  
redshift_vsolution/content/files/flowagent/redshift/
```

6. Create a properties file as `redshift_vsolution/content/files/flowagent/redshift.properties`. Configure the driver path relative to the location of the properties file as follows:

Sample Code

```
AMAZON_REDSHIFT_DRIVERMANAGER=./redshift/libamazonredshiftdbc64.so
```

7. Compress the vsolution in the `redshift_vsolution` directory:
`zip -r redshift_vsolution.zip ./`
8. Import the vsolution in System Management by performing the following substeps:

Note

You must have the Tenant Administrator role.

1. Start the System Management application from the Launchpad.
2. Open the *Tenant* tab.
3. Select **+** (*Add Solution*), and then select the `redshift_vsolution.zip` file that you created.
4. Open the *Strategy* tab, and then select **✎** (*Edit*).
5. Select **+** (*Add Solutions*).
6. Select the solution `vsolution_redshift-1.0.0` and select *Add*.
7. Select *Save*.

Note

For changes to take effect, restart the Flowagent application.

Operations

The Amazon Redshift connection type allows for the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Run rules in Metadata Explorer.
- Use as a remote source for rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Version	Amazon Redshift version.
Host	Host name of the Redshift server.
Port	Port number of the Redshift server.
Database name	Database name with which to connect.
SSL Mode	SSL certificate verification mode to use when connecting. Options include the following: <ul style="list-style-type: none">• prefer: if the server supports, the data is encrypted.• disable: data isn't encrypted.• allow: if server requires it, data is encrypted• require: data is always encrypted.• verify-ca: data is always encrypted, and server certificate is validated.• verify-full: data is always encrypted, server certificate is validated, and server hostname must match the one in the certificate. If you select verify-ca or verify-full, provide Redshift's certificate via the Certificates tab in Connection Management application.
User	Name of the Redshift user with privileges to connect to the database.
Password	Specified user's password.
Additional session parameters	Enter additional session-specific variables, if necessary.

Datatype Conversion

The application converts Redshift datatypes to an agnostic set of types, as shown in the following table.

Redshift Datatype	Datatype
BIGINT	BIGINT
DECIMAL(p, s)	DECIMAL(p, s)
DOUBLE PRECISION	FLOATING(8)
INTEGER	INTEGER(4)
REAL	FLOATING(4)
SMALLINT	INTEGER(4)
BOOLEAN	STRING(5)
CHAR(s)	STRING(s)
VARCHAR(s)	STRING(s)
DATE	DATE
TIMESTAMP	DATETIME
TIMESTAMPTZ	DATETIME
GEOMETRY	STRING(127)

Access Firewall

The network for Amazon Redshift instances is protected by a firewall that controls incoming traffic.

For SAP Data Intelligence cloud edition, Connection Management exposes an IP address using a read-only connection called `INFO_NAT_GATEWAY_IP`. Use this IP address and add the connection to allowlist in the Amazon Redshift dashboard. For SAP Data Intelligence on-premise edition, as a Kubernetes administrator, obtain the public IP address of all the nodes and add them to allowlist in the Amazon Redshift dashboard.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.35 RSERVER

The RServer connection type connects to and accesses information from an RServer server.

Operations

RServer is a TCP/IP server that allows other programs to communicate with one or several R sessions. You can create this connection type and use it in the R Client operator in the SAP Data Intelligence Modeler.

Attributes

Attribute	Description
Host	Host name or the IP address of the server.
Port	Port number of the server
User	User name to authenticate the RServer.
Password	Specified user's password for RServer.

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.36 S3

The Amazon Simple Storage Service (S3) connection type connects to and accesses information from objects in Amazon S3 or compatible services, such as Minio and Rook.

Operations

The S3 connection type allows for the following operations in Metadata Explorer:

- Browse remote objects.
- Obtain fact sheets of remote objects.
- Preview content of remote objects.
- Profile remote objects.
- Prepare data using the Preparation application.
- Save data from the Preparation application.

- Run rules.
- Extract metadata to catalog.

The S3 connection type allows for the following operations in the Modeler application:

- Read and write files using the Read File and Write File operators.
- Copy, rename, and remove files using the Copy File, Move File, and Remove File operators.

Attributes

Attribute	Description
Endpoint	<p>Endpoint URL of the Amazon S3 server. The protocol prefix isn't required. For example, <code>s3 . amazonaws . com</code>.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>When connecting to an Amazon S3-compatible service such as Minio, enter an endpoint that is an IP address rather than the hostname.</p> </div>
Protocol	<p>Values include the following:</p> <ul style="list-style-type: none"> • HTTP • HTTPS (default) <p>The value that you provide overwrites the value from the Endpoint attribute, if already set.</p>
Region	<p>Region over which you authenticate and operate. For example, <i>us-east-1</i>.</p>
Access Key	<p>ID of the key to access Amazon S3, in combination with the secret key.</p> <p>The Amazon Web Service (AWS) account must have permission to use the AWS S3 service. Also, some authorizations for specific services, such as S3, must be maintained in AWS identity and access management.</p>
Secret Key	<p>Secret access key to access AWS S3, in combination with the access key.</p> <p>The AWS account must have permission to use the AWS S3 service. Also, some authorizations for specific services, such as S3, must be maintained in AWS identity and access management.</p>

Attribute	Description
Root Path	<p>Optional root path name for browsing objects. The value must start with the forward slash character. For example, <code>/My Folder/MySubFolder</code>.</p> <p>If you specify the Root Path, then the application prefixes any path used with this connection with the Root Path.</p>
Use Assume Role	<p>Specifies whether to use temporary security credentials and restrict access to your S3 buckets based on an Identity and Access Management role (IAM role).</p> <p>Default is false.</p>
Role ARN	<p>Specifies the Amazon resource name (ARN) of the assumed IAM role. The Role ARN consists of your temporary access credentials.</p> <p>Required when you set <i>Use Assume Role</i> to true.</p> <div data-bbox="804 925 1394 1151" style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"> <p>Note</p> <p>Securely configure IAM Role by enforcing least privileges. For more information about security, see Security Recommendations for SAP Data Intelligence [page 272].</p> </div>
Role Session Name	<p>Specifies the name that uniquely identifies the assumed role session. Applicable only when the same Role ARN is used by others.</p>
Duration of Role Session	<p>Specifies the time in seconds for AWS to keep the session open. Default duration is 3600 seconds.</p>
External ID	<p>Specifies an assigned ID that ensures only specified third parties can assume the role.</p>
Role Policy	<p>Specifies the respective IAM policy in JSON format. Applicable when you use a session policy.</p>
Use Server Side Encryption	<p>Specifies whether to use S3 objects encrypted through server-side encryption.</p> <p>Default is false.</p>

Attribute	Description
Encryption Option	<p>Specifies the encryption to use. Options include the following:</p> <ul style="list-style-type: none"> • Encryption with Amazon S3 Managed Keys (SSE-S3): (default) specifies that your objects are encrypted using the default encryption configuration for objects in your Amazon S3 buckets. • Encryption with AWS Key Management Service Keys (SSE-KMS): specifies that your Amazon S3 buckets are configured to use AWS KMS keys. <p>For more information about each of these encryption options, see Server Side Encryption in your <i>Amazon Simple Storage Service User Guide</i>.</p>
KMS Key ARN	<p>Specifies the KMS key ARN for the customer-managed key, which AWS creates to encrypt the objects in your Amazon S3 buckets. The value must begin with 'arn'.</p> <p>Required when you set <i>User Server Side Encryption</i> to true and you choose Encryption with AWS Key Management Service Keys (SSE-KMS) for <i>Encryption Option</i>.</p>

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.37 SAP_IQ

The SAP_IQ connection type connects to and accesses information from objects in SAP IQ databases.

Operations

The SAP_IQ connection type allows for the following operations:

- View connection status in Connection Management.
- On-premise connectivity using SAP Cloud Connector (SCC).

The SAP_IQ connection type allows for the following operations in Metadata Explorer:

- Browse remote objects
- View fact sheets of remote objects
- Preview content of remote objects
- Profile remote objects

- Prepare data using the Preparation application
- Save data from the Preparation application
- Run rules
- Use as a remote source for rules
- Extract metadata to catalog
- Extract data lineage to catalog

The SAP_IQ connection type allows for the following operations in the Modeler application:

- Read tables and views
- Read data from SQL queries
- Process native SQL DDL/DML statements

Attributes

Attribute	Description
Version	SAP IQ version.
Host	SAP IQ server's host name.
Port	SAP IQ server's port number.
Use TLS	Specify whether to enable an encrypted connection.
Validate host certificate	Specify whether to validate the server certificate.
Hostname in certificate	Host name to validate in certificate when Validate host certificate is set to true (on).
Database name	Name of the database with which to connect.
Server name	SAP IQ server name with which to connect (also referred to as the Engine name).
User	SAP IQ user name that has privileges to connect to the specified database.
Password	Password of the specified SAP IQ user.
Additional session parameters	Specify additional session parameters, if necessary.

Datatype Conversion

The application converts SAP IQ datatypes to an agnostic set of types, called SAP Data Intelligence datatypes, as shown in the following table. The application supports only SAP IQ datatypes that have a corresponding SAP Data Intelligence datatype.

SAP IQ Datatype	SAP Data Intelligence Datatype
BIGINT	DECIMAL(20, 0)

SAP IQ Datatype	SAP Data Intelligence Datatype
UNSIGNED BIGINT	DECIMAL(20, 0)
BIT	INTEGER(4)
DECIMAL(p,s)	DECIMAL(p, s)
DOUBLE	FLOATING(8)
FLOAT	FLOATING(4)
INT	INTEGER(4)
UNSIGNED INT	INTEGER(4)
INTEGER	INTEGER(4)
UNSIGNED INTEGER	INTEGER(4)
MONEY	DECIMAL(19,4)
NUMERIC(p,s)	DECIMAL(p, s)
REAL	FLOATING(4)
SMALLINT	INTEGER(4)
SMALLMONEY	DECIMAL(10,4)
TINYINT	INTEGER(4)
UNIQUEIDENTIFIER	
UNIQUEIDENTIFIERSTR	STRING(36)
CHAR(s)	STRING(s)
VARCHAR(s)	STRING(s)
SYSNAME	STRING(30)
DATE	DATE
DATETIME	DATETIME
SMALLDATETIME	DATETIME
TIME	TIME
TIMESTAMP	DATETIME
BINARY	
VARBINARY	
BLOB	LARGE_BINARY_OBJECT
CLOB	LARGE_CHARACTER_OBJECT
IMAGE	LARGE_BINARY_OBJECT
LONG BINARY	LARGE_BINARY_OBJECT
TEXT	LARGE_CHARACTER_OBJECT

Note

For information about supported versions, see SAP Note [2693555](#).

4.4.38 SDL

The SDL (Semantic Data Lake) connection type connects to and accesses information from remote object stores.

Connection

Use the predefined connection `DI_DATA_LAKE` to access the SDL. This connection is managed by SAP and can't be modified.

Operations

An SDL connection allows for the following operations in Metadata Explorer:

- Browse remote objects.
- Obtain fact sheets of remote objects.
- Preview content of remote objects.
- Profile remote objects.
- Prepare data using the Preparation application.
- Save data from the Preparation application.
- Run rules.
- Extract metadata to catalog.

In Machine Learning scenarios, use an SDL connection for the following operations:

- Model train with artifact producer and artifact consumer.
- Model serving with artifact producer and artifact consumer.

In the Modeler application, use an SDL connection for the following operations:

- Read and write data using the Read File, Write File, and Flowagent File operators.
- Perform data operations using the Data Transform operator.

Attributes

Attribute	Description
ADLv2	
Authorization Method	Authorization method to use.
Account Name	Account name for the shared key authorization.
Account Key	Key for the shared key authorization.

Attribute	Description
Endpoint Suffix	If you don't enter an endpoint suffix, the system uses the default value <code>core.windows.net</code> .
WASB	
Account Name	Account name for WASB authorization.
Account Key	Key for the WASB authorization.
Endpoint Suffix	If you don't enter an endpoint suffix, the system uses the default <code>core.windows.net</code> .
GCS	
Project ID	ID of the GCP project.
Key File	Contents of the key file for authentication.
ADLv2, WASB, and GCS	
Root Path	<p>Root path name for browsing. The value starts with a forward slash and includes the bucket name. For example, <code>/MyBucket/My Folder</code>.</p> <p>Dataset names for this connection don't contain segments of the root path; instead their first segment is a subdirectory of the root path.</p>
Amazon S3	
Endpoint	Gateway endpoint to your S3 service. If you leave blank, the system uses the default value of <code>s3.amazonaws.com</code> .
Access Key ID	Access Key ID to authenticate your S3 connection.
Secret Access Key	Secret access key to authenticate your S3 connection.

Usage

The SDL connection includes the following directories:

- `/`: the root directory. You can't create directories or files under the root directory.
- `/shared/sap`: reserved for storing content and data produced by SAP components.
- `/external`: provides read-only access to connections for an authorized user.

Note

SAP doesn't support accessing other SDL connections through the `/external` directory.

- `/shared`: available by default for read and write access on the Object Store Type, which is defined in the connection.

Note

Previous releases included a `/work` directory intended for immutable artifacts. If your version includes a `/work` directory, it's deprecated and is now read-only. You can migrate the content from the `/work` directory to a new directory. After migration, a Tenant Administrator can delete the `/work` directory permanently.

4.4.39 SFTP

The SSH File Transfer Protocol (SFTP) connection type connects to and accesses information from an SSH File Transfer Protocol server.

Operations

The SFTP connection allows for the following operations:

- Read and write files using the Read File and Write File operators in the Modeler application.
- Remove files using the Remove File operator in the Modeler application.

Attributes

Attribute	Description
Hostname	Host name or IP address of the SFTP server.
Port	Port number of the SFTP server.
Host Key	The public key for the host, which is in the <code>.pub</code> format. For information about the process to retrieve a host public key, see Obtaining Host Key [page 195] . explains the process to retrieve a host's public key.
Authentication Type	Authentication type to use. Options include the following: <ul style="list-style-type: none">• SSH_Key (default)• password
User	Name of the user who is accessing the SFTP server.
Private Key	Specified user's SSH Private Key for Authentication Type SSH_Key. The server must already know the user's SSH Public Key.
Use Passphrase	Specify whether to use a passphrase for the private key.
Passphrase	Passphrase that decrypts the Private Key for Authentication Type SSH_Key.

Attribute	Description
Password	Specified user's password for Authentication Type password.
Root Path	Optional root path name for browsing objects. The value starts with the forward slash character. For example, /My Folder/MySubfolder. When you specify a Root Path, the application prefixes the specified root path to any path that you use with your WASB connection.
Proxy	Connection-specific proxy server (optional). Specify type, host, and port. SOCKS (default) type is available.

Host Key

The expected format of the file provided in the host key is one or more lines, each composed of the following elements:

```
<server host key algorithm> <SHA-256 fingerprint> <optional-comment>
```

❁ Example

The following is a valid file with two entries:

```
ssh-ed25519 AAAAC3NzaC1lZDI1NTE5AAAAIEMXBFYDfcYMW0dccgbJ/
TfhpTQhc5oR06jKIg+Wcarr myuser@myhost
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDRqWbaMxSetrsAtTHFaxym4rVqV1yb4umqhDJbJ0H63T+wn8
lm+Ev/i/
u+8BZT9nvzXZqbn1rezWZvXK234SkfDFzTib37vqlgPagrZlUc9DGAey6F4irQcgEQjiSaczsjNzYu
n2yrpsL/
9QBahFdeCKUPNQIXYU8ctbEOxqiOzvsNH4EsobiAS+leteRA0Pe2hiOaTODj4o3e5Pug4hugr8p/
tJPFVC5z7MBX9XPs6qpSAs8loZ0hZYdF4bjfHmaTNJTjrJcFg4RHTBVPsBKOLFBxwPhjcQlccNQ33v
oYF59bM37IyqV6h+Mz8up/
GrMVA7ka6np3fAyJhGhRPsLEZZY8h6KK633HLDqglkisQP87ewz8SRrcIHnhrP3hTBClx484XxCBmW
l4pUElQ+p32322v+KbwCEHpYj5pitnieekiXpsMNXOCZdyA/
llToPqzlGkbcI3z8ScOLvoX2qsrjOWMJlKOpwIcA/NzwU/
9LlFsecQvFzGowYYFHMnDypAnhCcwQz9BvqmjRRJGbmONmq39HTBmd0rfyoui8KCOGkN/
d89aZERzH6jZa9ft6qzaBuhKc1TND/ml+IBEUoWZUX3XurYaJu/
0awAcjVeyB0dhGafSRGhskBy2o0lX97Z0oErkIoc5BQRCLpa30jHywzd6BLnTKKJRS6pvfG9w==
```

Obtaining Host Key

Provide the host public key through a trusted channel. If your Windows 10, Linux, or MacOS machine has a trusted channel, perform the following steps by replacing the following elements with the specified values:

- \$HOST with the host name value of your connection
- \$PORT with the port value of your connection

Use the resulting file `host_key.pub.txt`, in the directory where you run the specified command, as the Host Key for your connection. The specified commands are as follows:

- **Windows 10:** In PowerShell, run the following command:

```
(ssh-keyscan -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 installed in your system), obtain the key through the following command:

```
ssh-keyscan -p $PORT $HOST 2>/dev/null | sed "s/^[^ ]* //" > host_key.pub.txt
```

4.4.40 SMTP

The Simple Mail Transfer Protocol (SMTP) connection type connects to and accesses information from an SMTP server connection.

Operations

The SMTP connection type allows for the following operations:

- Provide protocol to send e-mails.
- Use the protocol standard defined in RFC 2821.
- Create an SMTP connection for use in the Send Email operator in the Modeler.

Attributes

Attribute	Description
Host	Name of the host for the SMTP server.
Port	SMTP server port number.
User	User name to authenticate the SMTP server.
Password	Specified user's Password.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.41 SNOWFLAKE

The Snowflake connection type connects to and accesses information from Snowflake databases.

Prerequisites

Before you can use the Snowflake connection type, you must create a default Snowflake warehouse. Also, you must download and install the applicable ODBC driver. To download and install the ODBC driver, perform the steps in [Download and Install ODBC Driver for Snowflake \[page 199\]](#).

Operations

The Snowflake connection type allows the following operations:

- View connection status in Connection Management.
- Browse remote objects in Metadata Explorer.
- Obtain fact sheets of remote objects in Metadata Explorer.
- Preview content of remote objects in Metadata Explorer.
- Publish remote objects in Metadata Explorer.
- Run rules in Metadata Explorer.
- Extract metadata to catalog in Metadata Explorer.
- Extract data lineage to catalog in Metadata Explorer.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Process native SQL DDL/DML statements in Modeler.

Attributes

Attribute	Description
Host	Host name of the Snowflake server.
Port	Port number of the Snowflake server. The default value is 443.
Database	Name of the database with which to connect.
Warehouse	Name of the warehouse to use.

Attribute	Description
Authentication Type	Specifies the type of authentication for this connection. Options include the following: <ul style="list-style-type: none"> Password: (default) requires that you complete the <i>User</i> and <i>Password</i> attributes. KeyPair: requires that you complete either the <i>User</i> and <i>Private Key</i> attributes, or the <i>User</i>, <i>Private Key</i>, and <i>Passphrase</i> attributes.
Private Key	Browse to or enter the location of the private key file for KeyPair authentication.
Use Passphrase	Specifies whether to use a private key passphrase for KeyPair authentication.
Passphrase	Optional. Passphrase that decrypts the private key. You can create and use a substitution parameter for this attribute.
User	Name of a user with privileges to connect to the Snowflake database.
Password	Password for the named Snowflake user. Not applicable when you choose KeyPair <i>Authentication Type</i> .
Additional session parameters	Enter session-specific variables, if necessary.

Datatype Conversion

SAP Data Intelligence converts Snowflake datatypes to an agnostic set of types as described in the following table.

SAP HANA Data Lake Database Datatype	Mapped Datatype
BIGINT	DECIMAL(20, 0)
UNSIGNED BIGINT	DECIMAL(20, 0)
BIT	INTEGER(4)
DECIMAL(p,s)	DECIMAL(p, s)
DOUBLE	FLOATING(8)
FLOAT	FLOATING(4)
INT	INTEGER(4)
UNSIGNED INT	INTEGER(4)
INTEGER	INTEGER(4)

SAP HANA Data Lake Database Datatype	Mapped Datatype
UNSIGNED INTEGER	INTEGER(4)
MONEY	DECIMAL(19,4)
NUMERIC(p,s)	DECIMAL(p, s)
REAL	FLOATING(4)
SMALLINT	INTEGER(4)
SMALLMONEY	DECIMAL(10,4)
TINYINT	INTEGER(4)
UNIQUEIDENTIFIER	
VARCHAR(s)	STRING(s)
SYSNAME	STRING(30)
DATE	DATE
DATETIME	DATETIME
SMALLDATETIME	DATETIME
TIME	TIME
TIMESTAMP	DATETIME
BINARY	
VARBINARY	
BLOB	LARGE_BINARY_OBJECT
CLOB	LARGE_CHARACTER_OBJECT
IMAGE	LARGE_BINARY_OBJECT
LONG BINARY	LARGE_BINARY_OBJECT
TEXT	LARGE_CHARACTER_OBJECT

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.4.41.1 Download and Install ODBC Driver for Snowflake

A Snowflake connection requires that you first download and install the applicable ODBC driver.

Procedure

1. Download the Snowflake ODBC driver for Linux from <https://sfc-repo.snowflakecomputing.com/odbc/linux/index.html>. Use the following criteria:

- Select version 3.12.0.
 - Select the format TGZ (TAR file compressed using GZIP).
2. Create the vsolution area as `mkdir -p snowflake_vsolution/content/files/flowagent`.
 3. Create the vsolution manifest as `snowflake_vsolution/manifest.json`.

Sample Code

```
{
  "name": "vsolution_snowflake",
  "version": "<version.number>",
  "format": "2",
  "dependencies": []
}
```

Note

To upload a new driver later, modify the version to upload a new vsolution. Then modify your layering strategy appropriately.

4. Extract the downloaded compressed TAR archive, for example:

Example

```
tar -xvzf snowflake_linux_x8664_odbc-3.12.0.tgz -C snowflake_vsolution/
content/files/flowagent/
```

5. Create the following properties file with the driver manager relative to the location of the properties file: `snowflake_vsolution/content/files/flowagent/snowflake.properties` with driver manager relative to the location of the properties file:.

Sample Code

```
SNOWFLAKE_DRIVERMANAGER=./snowflake_odbc/lib/libSnowflake.so
PATH=./snowflake_odbc/lib:$PATH
LD_LIBRARY_PATH=./snowflake_odbc/lib:$LD_LIBRARY_PATH
```

Note

The exact path for the `SNOWFLAKE_DRIVERMANAGER` variable is based on the downloaded driver version.

6. Create a driver configuration file as follows: `snowflake_vsolution/content/files/flowagent/snowflake_odbc/lib/simba.snowflake.ini`. Use the following contents:

```
[Driver]
ANSIENCODING=UTF-8
DriverManagerEncoding=UTF-16
DriverLocale=en-US
ErrorMessagePath=./ErrorMessages
LogLevel=0
LogNamespace=
LogPath=
CURLVerboseMode=false
CABundleFile=./cacert.pem
ODBCInstLib=/usr/lib64/libodbcinst.so
```

7. Compress the vsolution from the `snowflake_vsolution` directory as follows:

```
cd snowflake_vsolution
zip -r snowflake_vsolution.zip ./
```

8. Import the vsolution in SAP Data Intelligence System Management by performing the following substeps:

Note

To perform these substeps, you must have the Tenant Administrator role.

- a. Open the System Management *Tenant* tab and choose **+** (*Add Solution*).
- b. Choose the `snowflake_vsolution.zip` file.
- c. Open the *Strategy* tab and choose **✎** (*Edit*).
- d. Choose **+** (*Add Solutions*).
- e. Choose the `vsolution_snowflake-<version>` solution and select *Add*.
- f. Select *Save*.
- g. Restart the Flowagent application.

Note

Changes take effect only after you restart the Flowagent application.

4.4.42 TERADATA

The Teradata connection type connects to and accesses information from Teradata databases.

Prerequisites

Before you create and use a Teradata connection type, perform the following steps:

1. Install the Teradata ODBC Driver for Linux Operating Systems. See your Teradata documentation for details.
2. Choose the linux version `x8664.17.10.00.14-1.tar.gz`, download the RPM file `tdodbc1710*.rpm`, and extract the RPM file.
3. Create the vsolution area as shown in the sample code:

Sample Code

```
mkdir -p teradata_vsolution/content/files/flowagent/teradata
```

4. Create the vsolution manifest as `teradata_vsolution/manifest.json` as shown in the sample code:

Sample Code

```
{
```

```
"name": "vsolution_teradata",
"version": "1.0.0",
"format": "2",
"dependencies": []
}
```

To upload a new driver later, modify the version, such as 1.0.1, 2.0.0, and so on, to upload a new vsolution. Then modify your layering strategy appropriately.

5. Extract the RPM file contents based on your platform:
 - On Windows, extract the RPM file using a file archiver.
 - On Linux, extract it using `rpm2cpio` tool:

↔ Sample Code

```
rpm2cpio tdodbc1710-17.10.00.14-1.x86_64.rpm |cpio -idmv
```

6. Copy extracted files to vsolution using the following sample code as a guide:

↔ Sample Code

```
cp -R tdodbc1710/opt/teradata/client/17.10/* teradata_vsolution/content/
files/flowagent/teradata/
```

7. Create a properties file: `teradata_vsolution/content/files/flowagent/teradata.properties`. Use a driver path relative to the location of the properties file, for example:

↔ Sample Code

```
TERADATA_17_DRIVERMANAGER=./teradata/tdataodbc_sb64.so
```

8. Compress the vsolution from within the `teradata_vsolution` directory, for example:

↔ Sample Code

```
cd teradata_vsolution
zip -r teradata_vsolution.zip ./
```

9. Import the vsolution in System Management:
 1. Start the System Management application from the Launchpad.

ⓘ Note

You must have the Tenant Administrator role.

2. Open the *Tenant* tab.
3. Select **+**, and select the `teradata_vsolution.zip` file that you created.
4. After the import is complete, open the *Strategy* tab, and select *Edit*.
5. Select **+**, and select the `vsolution_teradata-1.0.0` solution that you imported.
6. Select *Save*.
7. Restart the Flowagent application for the changes to take effect.

Operations

The Teradata connection type allows the following operations:

- View connection status in Connection Management.
- Read tables and views in Modeler.
- Read data from SQL queries in Modeler.
- Run native SQL DDL/DML statements in the Modeler and Metadata Explorer.

Attributes

Attribute	Description
Version	Teradata database version.
Host	Host name of the Teradata database server.
Port	Port number of the Teradata database server. The default value is 1025.
Database name	Name of the database with which to connect.
User	Teradata database user name with privileges to connect to the database.
Password	Password of the specified Teradata database user.
Additional session parameters	Enter any session-specific variables, if necessary.

Datatype Conversion

The application converts Teradata datatypes to an agnostic set of types, as shown in the following table. The application supports only Teradata datatypes with a corresponding datatype.

SAP HANA Data Lake Database Datatype	Mapped Datatype
BYTEINT	INTEGER(1)
SMALLINT	INTEGER(2)
INTEGER	INTEGER(4)
BIGINT	DECIMAL(28, 0)
NUMERIC	DECIMAL
DECIMAL(p, s)	DECIMAL(p, s)
NUMBER	DECIMAL
DOUBLE	FLOATING
FLOAT	FLOATING
REAL	FLOATING

SAP HANA Data Lake Database Datatype	Mapped Datatype
CHAR(s)	STRING(s)
VARCHAR(s)	STRING(s)
DATE	DATE
TIME	TIME
TIMESTAMP	DATETIME
BYTE	STRING(1)
VARBYTE(s)	STRING(s)
BLOB	LARGE_BINARY_OBJECT
CLOB	LARGE_CHARACTER_OBJECT

4.4.43 WASB

The Microsoft Windows Azure Storage Blob (WASB) connection type connects to and accesses information from objects in a WASB file system.

Operations

The WASB connection type allows the following operations in Metadata Explorer:

- Browse remote objects.
- Obtain fact sheets of remote objects.
- Preview content of remote objects.
- Profile remote objects.
- Prepare data using the Preparation application.
- Save data from the Preparation application.
- Run rules.
- Extract metadata to catalog.

The WASB connection allows the following operations in the Modeler application:

- Read and write files using the Read File and Write File operators.
- Copy, rename, and remove files using the Copy File, Move File, and Remove File operators.

Attributes

Attribute	Description
Protocol	Choose from the following options: <ul style="list-style-type: none">wasbs: uses HTTPS (default)wasb: uses HTTP
Account Name	Account name used in the Shared Key authentication.
Account Key	Account key used in the Shared Key authentication.
Endpoint Suffix	Default value is <code>core.windows.net</code> .
Root Path	Optional root path name for browsing objects. The value starts with the forward slash character. For example, <code>/MyFolder/MySubfolder</code> . When you specify a Root Path, the application prefixes the specified root path to any path that you use with your WASB connection.

Note

For more information about supported remote systems and data sources, see SAP Note [2693555](#).

4.5 Supported Connections in SAP Data Intelligence

SAP Data Intelligence supports data source and target systems and remote system orchestration.

For a list of all supported connections types in SAP Data Intelligence Connection Management and Modeler, see [Supported Connection Types \[page 122\]](#).

Related Information

[Supported Data Source Systems \[page 206\]](#)

[Supported Data Target Systems \[page 215\]](#)

[Remote System Orchestration \[page 218\]](#)

[Replication Flow Connections \[page 219\]](#)

4.5.1 Supported Data Source Systems

Many data source systems are supported in SAP Data Intelligence.

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
SSH File Transfer Protocol	n/a	SFTP [page 194]	n/a	Read File Write File Remove File List Files Monitor Files
Amazon Web Services Redshift	n/a	REDSHIFT [page 182]	n/a	Redshift Table Consumer Redshift SQL Consumer Table Consumer V2
Amazon Web Service Storage Service	n/a	S3 [page 186]	n/a	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files
Microsoft Azure Data Lake Storage Gen2	n/a	ADL_V2 [page 132]	n/a	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
Microsoft Azure Data Lake	n/a	ADL (Deprecated) [page 131]	n/a	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files
Windows Azure Storage Blobs	n/a	WASB [page 204]	n/a	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files
WampServer	n/a	WAMP	n/a	WAMP Consumer WAMP Producer
Amazon Web Services Simple Notification Service	n/a	AWS_SNS [page 134]	n/a	AWS SNS Consumer AWS SNS Producer
Simple Mail Transfer Protocol	n/a	SMTP [page 196]	n/a	Send Email Notification
Alibaba Cloud Object Storage Service	n/a	OSS [page 179]	n/a	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
Azure SQL Database	n/a	AZURE_SQL_DB [page 135]	n/a	Azure SQL DB SQL Consumer (Deprecated) Azure SQL DB Table Consumer (Deprecated) Table Consumer Table Replicator V3 (Deprecated)
IBM DB2	10.x and above	DB2 [page 143]	n/a	DB2 SQL Consumer (Deprecated) DB2 Table Consumer (Deprecated) Table Consumer Table Replicator V3 (Deprecated)
Google BigQuery	n/a	GCP_BIGQUERY [page 145]	Key File	Google BigQuery SQL Consumer (Deprecated) Table Consumer
Google Cloud Pub/Sub	n/a	GCP_PUBSUB [page 149]	Key File	Google Pub/Sub Consumer Google Pub/Sub Producer
Google Cloud Storage	n/a	GCS [page 149]	Key File	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
Hadoop Distributed File System	n/a	HDFS [page 154]	n/a	Structured File Consumer V3 Structured File Producer V3 Read File Write File Remove File List Files Monitor Files
Hypertext Transfer Protocol (Secure)	n/a	HTTP [page 160]	n/a	HTTP Client HTTP Server OpenAPI Client OpenAPI Server Rest API Client
Internet Message Access Protocol	n/a	IMAP [page 161]	n/a	Receive Email
Kafka	Greater than or equal to version 0.8	KAFKA [page 163]	n/a	Kafka Consumer V1 Kafka Producer V1
Microsoft SQL Server (on-prem)	2012, 2014, 2016, 2017	MSSQL [page 169]	n/a	SQL Server SQL Consumer (Deprecated) SQL Server Table Consumer (Deprecated) Table Consumer Table Replicator V3 (Deprecated)
Microsoft SQL Server on Azure VM (IaaS)	2012, 2014, 2016, 2017	MSSQL [page 169]	n/a	SQL Server SQL Consumer (Deprecated) SQL Server Table Consumer (Deprecated) Table Consumer Table Replicator V3 (Deprecated)
MQTT	n/a	MQTT	n/a	MQTT Consumer MQTT Producer

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
MySQL	5.5, 5.6	MYSQL [page 170]	n/a	MySQL SQL Consumer (Deprecated) MySQL Table Consumer (Deprecated) Table Consumer Table Replicator V3 (Deprecated)
NATS	n/a	NATS	n/a	NATS Consumer NATS Producer
OData	V2, V4	ODATA [page 171]	oAuth2	OData Query Consumer SAP Application Consumer SAP Application Producer
Oracle	10g, 11g, 12c, 18c, 19c	ORACLE [page 176]	basic (user+pw)	Oracle SQL Consumer (Deprecated) Oracle Table Consumer (Deprecated) Table Consumer Table Replicator V3 (Deprecated)
R Server	R 3.3.3 und 3.5.1	RSERVE [page 186]	n/a	R Client (Deprecated)
SAP Business Warehouse (BW)	7.40 SP08 and above 7.50 SP04 and above 7.51 SP00 and above 7.52 SP00 and above	BW [page 137]	n/a	Data Transfer
SAP Business Warehouse (BW)	7.30 SP14 and above with ODP API 2.0	ABAP LEGACY [page 127]	n/a	SAP ABAP ODP Object Consumer (Deprecated) SAP Application Consumer (full load only)
SAP Business Warehouse (BW)	7.30 SP20 and above with ODP API 2.0	ABAP [page 124]	n/a	ABAP ODP Reader

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
SAP BW/4 HANA	1.0 SP04 and above	BW [page 137]	n/a	Data Transfer BW Process Chain SAP Application Producer (via aDSO Write Interface)
SAP S/4 HANA Cloud Edition	n/a	ABAP [page 124]	n/a	ABAP CDS Reader
SAP S/4 HANA on-premise	1909 and above	ABAP [page 124]	n/a	Data Transfer ABAP ODP Reader SAP ABAP Operator SLT Connector Custom ABAP Operator SAP Application Producer (via OData)
SAP S/4 HANA on-premise	>= 1610 and < 1909 with dedicated SLT Server	ABAP [page 124]	n/a	SLT Connector Custom ABAP Operator SAP ABAP Operator
SAP S/4 HANA on-premise	any version that supports ODP API 2.0	ABAP LEGACY [page 127]	n/a	SAP ABAP ODP Object Consumer (Deprecated) SAP Application Consumer (full load only)
SAP S/4 Foundation	>= 105 with DMIS 2020, ODP with ODP API 2.0 and Minimum BW	ABAP [page 124]	n/a	SAP ABAP Operator SLT Connector Custom ABAP Operator ABAP ODP Reader
SAP NetWeaver (ODP-based extraction)	7.30 SP20 and above with ODP API 2.0 and Minimum BW	ABAP [page 124]	n/a	ABAP ODP Reader

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
SAP NetWeaver (Table / view based extraction)	7.00 or higher	ABAP [page 124]	n/a	SAP ABAP Operator Custom ABAP Operator SLT Connector Cluster Table Splitter (deprecated)
SAP NetWeaver (Table / view based extraction)	6.20 via dedicated SLT Server	ABAP [page 124]	n/a	SLT Connector
SAP NetWeaver	7.30 SP 14 and above with ODP API 2.0	ABAP LEGACY [page 127]	n/a	SAP ABAP ODP Object Consumer (Deprecated) SAP Application Consumer (full load only)
SAP Cloud Data Integration	n/a	CLOUD_DATA_INTEGRATION [page 138]	n/a	Cloud Data Integration Consumer
SAP BTP Open Connectors	n/a	OPEN_CONNECTORS [page 173]	n/a	Open Connectors SQL Consumer Open Connectors Table Consumer (Deprecated)
SAP IQ Server	16.x	SAP_IQ [page 189]	n/a	SAP IQ SQL Consumer (Deprecated) SAP IQ Table Consumer (Deprecated) Table Consumer V2 Flowagent Table Producer (Deprecated) Table Producer V3

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
SAP HANA	Edition 2.0 Edition 1.0 SPS12 and above	HANA_DB [page 150] (query HANA tables/ HANA views)	n/a	SAP HANA Client SAP HANA Monitor (Deprecated) HANA Table Consumer (Deprecated) Table Consumer V2 Flowagent Table Producer (Deprecated) Table Producer V3 SQL Producer Write to HANA Table Read HANA Table Run HANA SQL Initialize HANA Table HANA ML Forecast HANA ML Training HANA ML Inference

Connection	Version	Connection Type	Authentication	Associated SAP Data Intelligence Operators
SAP HANA Cloud	n/a	HANA_DB [page 150]	n/a	SAP HANA Client SAP HANA Monitor (Deprecated) HANA Table Consumer (Deprecated) Table Consumer V2 Flowagent Table Producer (Deprecated) Table Producer V3 SQL Producer Write to HANA Table Read HANA Table Run HANA SQL Initialize HANA Table HANA ML Forecast HANA ML Training HANA ML Inference
SAP HANA Data Lake Database	n/a	HDL_DB [page 156]	n/a	Table Consumer V2 Table Producer V3
SAP BTP Enterprise Messaging	n/a	CPEM [page 140]	n/a	SAP CP EM Consumer SAP CP EM Producer
Snowflake	n/a	SNOWFLAKE [page 197]	n/a	Table Consumer V2 SQL Consumer V2
Teradata	17	TERADATA [page 201]	Basic authentication	Table Consumer V2SQL Consumer V2 Flowagent SQL Executor V2 Flowagent SQL Executor V1 (Deprecated)

4.5.2 Supported Data Target Systems

SAP Data Intelligence supports many data target systems.

Note

The **Associated SAP Data Intelligence Operators** column includes only Generation 1 (Gen 1) operators.

Connection	Version	Connection Type	Associated SAP Data Intelligence Operators (Gen 1 only)
SSH File Transfer Protocol	n/a	SFTP [page 194]	Binary File Producer Remove File
Amazon Web Service Storage Service	n/a	S3 [page 186]	Structured File Producer V3 Binary File Producer Remove File
Microsoft Azure Data Lake Storage Gen2	n/a	ADL_V2 [page 132]	Structured File Producer V3 Binary File Producer Remove File
Microsoft Azure Data Lake	n/a	ADL (Deprecated) [page 131]	Structured File Producer V3 Binary File Producer Remove File
Windows Azure Storage Blobs	n/a	WASB [page 204]	Structured File Producer V3 Binary File Producer Remove File
WampServer	n/a	WAMP	WAMP Producer
Amazon Web Services Simple Notification Service	n/a	AWS_SNS [page 134]	AWS SNS Producer
Alibaba Cloud Object Storage Service	n/a	OSS [page 179]	Structured File Producer V3 Binary File Producer Remove File
Google BigQuery	n/a	GCP_BIGQUERY [page 145]	Google BigQuery Table Loader (Deprecated) Cloud Table Producer
Google Cloud Pub/Sub	n/a	GCP_PUBSUB [page 149]	Google Pub/Sub Producer
Google Cloud Storage	n/a	GCS [page 149]	Structured File Producer V3 Binary File Producer Remove File

Connection	Version	Connection Type	Associated SAP Data Intelligence Operators (Gen 1 only)
Hadoop Distributed File System	n/a	HDFS [page 154]	Structured File Producer V3 Binary File Producer Remove File
Hypertext Transfer Protocol (Secure)	n/a	HTTP [page 160]	HTTP Client OpenAPI Client OpenAPI Server
Kafka	greater than or equal to version 0.8	KAFKA [page 163]	Kafka Producer V1
MQTT	n/a	MQTT	MQTT Producer
NATS	n/a	NATS	NATS Producer
OData	V2, V4	ODATA [page 171]	SAP Application Producer V2
SAP BW/4 HANA	1.0 SP04 and above	BW [page 137]	SAP Application Producer V2
SAP S/4 HANA on-premise	1909 and above	ABAP [page 124]	Custom ABAP Operator Writing back to S/4 HANA on Premise via customizing using Custom ABAP Operator
SAP S/4 Foundation	>= 105 with DMIS 2020	ABAP [page 124]	Custom ABAP Operator Writing back to S/4 Foundation via customizing using Custom ABAP Operator
SAP NetWeaver	7.00 or higher	ABAP [page 124]	Custom ABAP Operator Writing back to SAP NetWeaver systems via customizing using Custom ABAP Operator
SAP IQ Server	16.x	SAP_IQ [page 189]	Flowagent Table Producer (Deprecated) Table Producer V3

Connection	Version	Connection Type	Associated SAP Data Intelligence Operators (Gen 1 only)
SAP HANA	Edition 2.0	HANA_DB [page 150]	SAP HANA Client
	Edition 1.0 SPS12 and above		Flowagent Table Producer (Deprecated) Table Producer V3 SQL Producer Write to HANA Table Run HANA SQL Initialize HANA Table
SAP HANA Cloud	n/a	HANA_DB [page 150]	SAP HANA Client
			Flowagent Table Producer (Deprecated) Table Producer V3 SQL Producer Write to HANA Table Run HANA SQL Initialize HANA Table
SAP HANA Data Lake Database	n/a	HDL_DB [page 156]	Table Producer V3
SAP BTP Enterprise Messaging	n/a	CPEM [page 140]	SAP CP EM Producer
SAP HANA Data Lake Files	n/a	HDL_FILES [page 159]	Structured File Producer V3
			Binary File Producer Remove File
Snowflake	n/a		Cloud Table Producer

4.5.3 Remote System Orchestration

The following table lists the remote system orchestration supported in SAP Data Intelligence.

Connection	Version	Connection Type	Associated SAP Data Intelligence Operators*
SAP Data Services	v4.2 SP8 Patch 4 and above v4.2 SP9 Patch 1 and above v4.2 SP10 v4.2 SP11	DATASERVICES [page 142]	SAP Data Services Job
SAP Information Steward	4.2 SP14	INFORMATION_STEWARD [page 162]	n/a Connection Type offers ability to import rules and rule bindings in the Metadata Explorer
SAP Business Warehouse (BW)	7.40 SP08 and above 7.50 SP04 and above 7.51 SP00 and above 7.52 SP00 and above	BW [page 137]	BW Process Chain
SAP BW/4 HANA	1.0 SP04 and above	BW [page 137]	BW Process Chain
SAP HANA	Edition 2.0 Edition 1.0 SPS12 and above	HDL_DB [page 156]	HANA ML Forecast HANA ML Training HANA ML Inference
SAP HANA	Edition 2.0 Edition 1.0 SPS12 and above	HANA_XS [page 153] (remotely orchestrate XSC-based SDI Flowgraphs)	HANA Flowgraph (Deprecated)
SAP HANA Cloud	n/a	HANA_DB [page 150]	HANA ML Forecast HANA ML Training HANA ML Inference
SAP Data Intelligence	n/a	n/a	Pipeline
SAP BTP Integration System	n/a	CPI [page 141]	SAP CPI-PI iFlow
Google Cloud Dataproc	n/a	GCP_DATAPROC [page 148]	Submit Hadoop Job (Deprecated)

* Currently, the *Associated SAP Data Intelligence Operators* column in this table includes only Generation 1 operators.

4.5.4 Replication Flow Connections

The following tables list the data source and target systems supported by the replication management service in SAP Data Intelligence. This service manages the replication functionality (replication flows) in the SAP Data Intelligence Modeler application.

Prerequisites for Setting Up Connections for an SDC HANA Target Table

To write data to a target HANA database table, the following privileges are required for the HANA database connection user.

Privilege	Objects	Description
<ul style="list-style-type: none">• INSERT• UPDATE• DELETE	Target table	To apply DML (Data Manipulations Language) operations to the target table; truncate the target table if requested in a replication task.
SELECT	System views: SYS.TABLES, SYS.TABLE_COLUMNS, SYS.INDEX_COLUMNS	To query metadata of the existing target table.
CREATE ANY	Database schema of the target table (container in the RMS target space).	To create a target table by CREATE TABLE DDL if none exists.

Source Systems for SAP Data Intelligence

Connection	Connection Type	Notes
SAP HANA Cloud	HANA_DB [page 150]	Replication from tables stored in SAP HANA Cloud.
SAP S/4HANA Cloud	ABAP [page 124]	CDS view replication for version 2202 and later. Prerequisites for CDS view extraction: Loading and Replicating Data from ABAP CDS Views in SAP S/4HANA .

Connection	Connection Type	Notes
SAP S/4HANA on premise	ABAP [page 124]	<p>CDS view replication for version 1909 and later; TCI note implementation required:</p> <ul style="list-style-type: none"> • • SAP S/4HANA 1909 with TCI Note 3373487 and central note 2830276. • SAP S/4HANA 2020 with TCI Note 3373436 and central note 2943599. • SAP S/4HANA 2021 with TCI Note 3581682 and central note 3085579. • SAP S/4HANA 2022 with TCI Note 3581682 and central note 3130827. • SAP S/4HANA 2023 with TCI Note 3581682 and central note 3254953. <p>Prerequisites for CDS view extraction: Loading and Replicating Data from ABAP CDS Views in SAP S/4HANA.</p> <p>Additionally, table-based replication as well as ODP-based extraction is supported. Check the following Note for details: 2890171.</p> <div data-bbox="874 1061 1394 1227" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <p>ODP-based replication is supported for ODP context ODP_SAPI and ODP_BW, and requires ODP API version 2.</p> </div>

Connection	Connection Type	Notes
SAP ECC/SLT	ABAP [page 124]	<p>Table-based replication for DMIS 2018 SP06 or later with TCI note 3110660. See also the central note 3085541.</p> <p>Table-based replication for DMIS 2020 SP03 or later.</p> <p>ODP-based replication for DMIS 2011 SP23, DMIS 2018 SP08, DMIS 2020 SP04 or a high-level SP. Check the following Note for details: 2890171.</p> <div data-bbox="874 651 1398 1014"> <p>Note</p> <p>Older SAP systems, for example, those using DMIS 2011, can be integrated using a dedicated SLT system that meets the minimum prerequisites of DMIS 2018 SP06 or later (with TCI note 3110660 and note 3085541) or DMIS 2020 SP03 or later, and considers SLT version dependency. See the <i>SAP Landscape Transformation Replication Server Installation Guide</i>.</p> </div> <div data-bbox="874 1032 1398 1339"> <p>Note</p> <p>Replications from ABAP sources guarantee only eventual consistency in the target system.</p> <p>For replications of cluster tables via SLT, this can lead to temporarily deleted records in the target dataset when the source dataset is modified (inserts, modifications, deletes).</p> </div>
SAP Business Warehouse	ABAP [page 124]	<p>ODP-based replication for various SAP Business Warehouse data sets using DMIS 2011 SP23, DMIS 2018 SP08, DMIS 2020 SP04 or a higher level SP. Check the following Note for details: 2890171.</p> <div data-bbox="874 1529 1398 1697"> <p>Note</p> <p>ODP-based replication is supported for ODP context ODP_SAPI and ODP_BW and requires ODP API version 2.</p> </div>
Microsoft Azure SQL database	AZURE_SQL_DB [page 135]	Table-based replication from Microsoft Azure SQL databases.

Target Systems for SAP Data Intelligence

Connection	Connection Type	Notes
SAP HANA Cloud	HANA_DB [page 150]	Replication into tables stored in SAP HANA Cloud.
Amazon S3	S3 [page 186]	Replication into files stored in Amazon S3 using CSV, Apache Parquet, JSON, or JSONLines file formats.
Microsoft Azure Data Lake Storage Gen2	ADL_V2 [page 132]	Replication into files stored in Azure Data Lake Storage Gen2 using CSV, Apache Parquet, JSON, or JSONLines file formats.
Google Cloud Storage	GCS [page 149]	Replication into files stored in Google Cloud Storage using CSV, Apache Parquet, JSON, or JSONLines file formats.
HANA Data Lake Storage	HDL_FILES [page 159]	Replication into files stored in HANA Data Lake Storage using CSV, Apache Parquet, JSON, or JSONLines file formats.
Apache Kafka	KAFKA [page 163]	Replication into topics stored in Kafka using JSON or Avro serializations.

Note

Kafka version 0.8 and higher.

Note

If your connection configuration specifies paths to the SAP Data Intelligence repository (for example, the paths specified in *Certificate Authority Path*), the following steps are required:

1. Log in using the credentials for the user that uploaded the corresponding files to the workspace.
2. Open the System Management application and then open the *Files* tab.
3. Select one or more files, select **•••** (*Show actions*), and choose *Export as solution to solution repository*.
4. Enter a name and version and select *Export as solution*.
5. Log in as a tenant administrator.
6. Open the System Management application.
7. Open the *Tenant* tab and then open the *Strategy* tab.
8. Click *Edit* to add the exported solution. For more information, see [Manage Tenant \[page 78\]](#).

Related Information

[Replicating Data](#)

[SAP Note 2890171 - SAP Data Intelligence - ABAP Integration](#)

[SAP Note 2187425 - Information about SAP Note Transport based Correction Instructions \(TCI\)](#)

[SAP Landscape Transformation Replication Server Installation Guide](#)

4.6 Using SAP Cloud Connector Gateway

The SAP Cloud Connector is a proxy application that allows SAP Data Intelligence in the cloud to access on-premise connections.

Several connection types support using the SAP Cloud Connector gateway. If a connection type supports the gateway and the gateway is registered during the installation of the SAP Data Intelligence instance, there's an additional field named *Gateway* in the Create Connection or Update Connection page in the Connection Management application.

Note

Choose the SAP Cloud Connector gateway for your connection for the following connection types:

- ABAP (only RFC)
- ABAP_LEGACY
- BW
- CLOUD_DATA_INTEGRATION
- DATASERVICES
- DB2
- HANA_DB
- HTTP
- INFORMATION_STEWARD
- KAFKA

When you connect to KAFKA with SAP Cloud Connector, and the virtual host differs from the actual hostname, set the virtual hostname in the 'kafka_advertised_listeners' server configuration.

- MYSQL
- ODATA
- OPENAPI
- ORACLE
- MSSQL
- SAP_IQ
- SFTP
- Teradata

To establish connectivity to the on-premise connection using the SAP Cloud Connector proxy, choose *SAP Cloud Connector* in the *Gateway* field.

Except for the HTTP connection, you can access cloud connector instances with locations other than the default ID. Set the *Location ID* field in the *Gateway* section with a different ID.

Note

HTTP connection type uses only the default ID; it does **not** support Location ID.

Related Information

[Troubleshooting SAP Cloud Connector \[page 284\]](#)

[Configure Cloud Connector \[page 22\]](#)

4.7 (Mandatory) Configure Authorizations for Supported Connection Types

To perform the supported operations of a connection type, various privileges are required in source system.

You must configure minimum or mandatory authorizations that are needed at the source to perform all the supported operations and access data from SAP Data Intelligence.

Example

When you create a connection to the BW system and define the parameters to use the connection, you must also configure other authorizations in the BW system so that you can perform the supported operations of BW connection type.

Related Information

[HANA_DB \[page 224\]](#)

4.7.1 HANA_DB

Prerequisites to configure HANA database as a data source in SAP Data Intelligence.

Read or Extract Metadata

- HANA TABLES or SQL VIEWS
GRANT SELECT ON SCHEMA <specific_schema> to <connection_user>
GRANT SELECT ON "<specific_schema>". "<specific_tablename>" to <connection_user>

- CALCVIEWS (OLAP views are excluded)

```
GRANT SELECT ON SCHEMA <calcview_schema> to <connection_user>
GRANT SELECT ON "<calcview_schema >". "<specific_calcviewname>" to
<connection_user>
GRANT SELECT ON "_SYS_BI"."BIMC_DIMENSION_VIEW" to <connection_user>
GRANT SELECT ON "_SYS_BI"."BIMC_VARIABLE_VIEW" to <connection_user>
GRANT SELECT ON "_SYS_BI"."BIMC_VARIABLE_RANGE_DEFAULTS" to <connection_user>
GRANT SELECT ON "_SYS_BI"."BIMC_ALL_AUTHORIZED_CUBES" to <connection_user>
GRANT SELECT ON "_SYS_BI"."BIMC_DIMENSIONS" to <connection_user>
```

Change Data Capture and Load to HANA Table

```
GRANT ALTER|CREATE ANY|CREATE TEMPORARY TABLE|CREATE VIRTUAL PACKAGE|
DELETE|DROP|EXECUTE|INDEX|INSERT|TRIGGER|UPDATE ON SCHEMA <specific_schema> to
<connection_user>
```

```
GRANT ALTER|DELETE|DROP|EXECUTE|INDEX|INSERT|TRIGGER|UPDATE|REFERENCES ON
"<specific_schema>". "< specific_tablename>" to <connection_user>
```

4.8 Allowing SAP Data Intelligence Access Through Firewalls

Configure SAP Data Intelligence to access external data sources that are protected by a network access tool, such as a firewall.

To allow SAP Data Intelligence access to an external data source that is protected by a network access tool, add the SAP Data Intelligence network address translation (NAT) gateway IP address to the external data source allow-list. The NAT IP address is static and doesn't change during the lifetime of the SAP Data Intelligence cluster.

An SAP Data Intelligence cluster uses an NAT gateway to connect to the internet and other network endpoints that are available on its network, including endpoints accessed through the following methods:

- Virtual Private Cloud (VPC) peering
- VNet peering
- Virtual Private Network (VPN)

To find the NAT IP address, perform the following steps:

1. Open the Connection Management application.
2. Search for `INFO_NAT_GATEWAY_IP`.
3. Hover your mouse over the text in the **Description** column of the `INFO_NAT_GATEWAY_IP` connection.
4. The NAT Gateway IP address appears in the popup text.

Note

`INFO_NAT_GATEWAY_IP` is also available for clusters on Azure.

To find the Azure VNet Subnet ID, perform the following steps:

1. Open the Connection Management application.
2. Search for `INFO_VNET_SUBNET_ID`.
3. Hover your mouse over the text in the **Description** column of the `INFO_VNET_SUBNET_ID` connection.
4. The VNet Subnet ID appears in the popup text.

5 Monitoring SAP Data Intelligence

SAP Data Intelligence provides a stand-alone monitoring application to monitor the status of graphs run in the Modeler. The Monitoring application provides capabilities to visualize the summary of graphs run in the SAP Data Intelligence Modeler with relevant charts.

The Monitoring application also allows you to schedule graph runs. For each graph instance, the Monitoring application provides details for processes like the following:

- Graph run status.
- Time of graph run.
- Graph type.
- Graph source.

The Monitoring application allows you to open a graph in the Modeler, view graph configurations, or stop process runs.

You can also view the execution and configuration of replication flows and their associated tasks.

Monitoring Policy

Developer member users assigned the `sap.dh.monitoring` policy can view analytics and instances of graphs for all tenant users. However, the policy doesn't provide member user access to schedules.

Without the policy, member users can monitor only their own graphs.

Related Information

[Monitoring Graphs \[page 227\]](#)

[Log in to SAP Data Intelligence Monitoring \[page 229\]](#)

[Using the Monitoring Application \[page 229\]](#)

[Stop, Restart, and Pause Tenant User Graphs \[page 237\]](#)

[Accessing the SAP Data Intelligence Monitoring Query API \[page 238\]](#)

[Pre-Delivered Policies \[page 55\]](#)

5.1 Monitoring Graphs

After creating and running graphs, monitor graphs and view statistics.

The following table describes the options for monitoring the graph status in the SAP Data Intelligence Modeler.

Action	Description
Monitor Status of Graph Runs	<p>After you create and run a graph, monitor the status of the graph run in the Modeler.</p> <p>Use the standalone monitoring application that SAP Data Intelligence provides to monitor the status of all graphs executed in the Modeler.</p>
Trace Messages	<p>Trace messages monitor both the system and running graphs to isolate problems or errors that may occur. Trace messages provide an initial analysis of your running graphs so you can troubleshoot potential problems or errors.</p>
Use SAP Data Intelligence Monitoring	<p>SAP Data Intelligence provides the Monitoring application to monitor the status of graphs run in the SAP Data Intelligence Modeler.</p>
Access the SAP Data Intelligence Monitoring Query API	<p>Access the SAP Data Intelligence Monitoring Query API to retrieve application performance metrics for your tenant. For more information, see Accessing the SAP Data Intelligence Monitoring Query API [page 238].</p>
SAP Data Intelligence Diagnostics	<p>SAP Data Intelligence Diagnostics deploys one of the most widely used stacks of open-source monitoring and diagnostic tools for Kubernetes. For health and performance monitoring, SAP Data Intelligence Diagnostics provides cluster administrators access to cluster-wide system and application metrics.</p>

ⓘ Note

Developer member users with the `sap.dh.monitoring` policy can monitor the status of graph runs for all tenant users.

ⓘ Note

By default, the Modeler also performs logging. The log messages are intended for a broader audience with different skills. If you want to view the log messages, start the Modeler, and in the bottom pane, select the [Logs](#) tab.

5.2 Log in to SAP Data Intelligence Monitoring

Access the SAP Data Intelligence Monitoring application from the SAP Data Intelligence Launchpad or directly launch the application with a stable URL.

Prerequisites

You must have administrator permission to perform this task.

Procedure

1. Open the SAP Data Intelligence Launchpad URL in a browser.
The welcome screen opens.
2. Log into the SAP Data Intelligence Launchpad using the following information:
 - Tenant name
 - Username
 - Password

For new instance, the tenant name is "default". For user name and password, use the credentials that you used when you created a service.

Note

If you enter an incorrect password five consecutive times within a minute, your account is temporarily locked. Wait 10 seconds until your next attempt.

The SAP Data Intelligence Launchpad opens and displays the initial home page. To view user details, such as the tenant ID, select the profile icon in the upper right of the screen.

The home page displays the application tiles available in the tenant based on your assigned policies.

3. Select the *Monitoring* tile.

The Monitoring application opens to the *Analysis* tab.

5.3 Using the Monitoring Application

The SAP Data Intelligence Monitoring application offers the following capabilities.

Tenant administrator users, and developer member users with the `sap.dh.monitoring` policy, can view all tenant users statistics. Member users without the monitoring policy can view information only for their graphs.

Note

Developer member users are those who have the `sap.dh.developer` and `sap.dh.member` policies assigned to them.

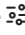


Developer member users that also have the `sap.dh.monitoring` policy can filter the information based on specific users.

Actions	Description
Visualize various aspects of graph execution in the <i>Analytics</i> tab.	The <i>Analytics</i> tab is a dashboard that contains various tiles with targeted information about graphs. The following table describes each tile of the <i>Analytics</i> tab.

Tile	Description
<i>Status</i>	<p>A pie chart with the following information:</p> <ul style="list-style-type: none"> • The number of graph instances executed in the Modeler. • The status of graph instances executed. Each sector in the pie chart represents a graph state.
<i>Runtime Analysis</i>	<p>A scatter chart that includes a time period axis and duration of graph execution (in seconds) axis. This chart provides information on each graph execution and plots them in the scatter chart against the time and duration of execution.</p> <p>Each point in the chart represents a graph instance. Place your cursor on one instance point for more information about that particular graph instance.</p> <p>Configure the chart to view results for a selected time period. Use the filter in the chart header to select the time period (hour, day, and week).</p>
<i>Recently Executed</i>	<p>Displays the top five instances and execution status. To view all the instances, select <i>Show All</i>.</p>
<i>Memory Usage</i>	<p>A line chart for the memory consumption of graphs. Use the filter to display by graph, status, and submission time. You can also see the resource usage in the last hour, day, 2 days, or set the custom time range for which to view the resource consumption.</p>
<i>CPU Usage</i>	<p>A line chart for the CPU consumption of graphs. Filter to display by graph, status, user, or submission time. Also see the resource usage in the last hour, day, 2 days, or set the custom time range for which you want to view the resource consumption.</p>

Actions	Description
	<p>Note</p> <p>The tiles in the <i>Analytics</i> tab don't include the information on the archived graph instances.</p>
<p>View execution details of individual graph instances in the <i>Instances</i> tab.</p>	<p>Use the <i>Instances</i> tab of the SAP Data Intelligence Monitoring application to view execution details of individual graph instances.</p> <p>Note</p> <p>If you're a tenant administrator, you can see graph instances of all users. By default, there's a filter that shows only the administrator's graph instances. Modify the filter to view graph instances of select tenant users on which you can perform limited actions.</p> <p>Note</p> <p>If you're a developer member user with the <code>sap.dh.monitoring</code> policy, you can see graph instances of all users, just like a tenant administrator can.</p> <p>For each graph instance, the Monitoring application provides information, such as the status of the graph execution, the graph execution name, the source of the graph in the repository, and the time of execution.</p> <p>To open the execution details pane, select a graph instance. The execution details pane displays more execution details that help you monitor the graph execution.</p> <p>Note</p> <p>By default, a filter is applied to exclude the subgraphs and the archived instances from the <i>Instances</i> list. You can remove the filter by selecting the × icon that corresponds to a filter in the <i>Filters</i> bar. Alternately, select the Filter icon to set or remove filters.</p>

Actions	Description
View subgraph execution details in the <i>Instances</i> tab.	<p>When a graph execution triggers or spawns the execution of another graph, the spawned execution is called a subgraph.</p> <p>You can filter the view of the Monitoring application to view execution details of all subgraph instances.</p> <p>To view subgraphs, remove the filter <i>Exclude Subgraphs</i> by selecting the X Delete icon that corresponds to a filter in the <i>Filters</i> bar. Or use the Filter icon to set or remove filters. The application refreshes the list view and displays all subgraph instances along with all other graph instances.</p> <div data-bbox="603 667 1394 786" style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <p>All subgraph instances are named as <i>subgraph</i>.</p> </div> <p>To open the execution details pane, select a subgraph instance. The subgraph instance pane displays additional execution details.</p>
View subgraph hierarchy details in the <i>Instances</i> tab.	<p>View the parent graph of a selected subgraph, or view the complete hierarchy of graph instances associated with the subgraph.</p> <ol style="list-style-type: none"> 1. Open the <i>Instances</i> tab. 2. Select the ... <i>Select an Action</i> icon next to the applicable subgraph instance. 3. Select <i>Show Hierarchy</i>. The Hierarchy dialog box opens displaying all graph instances associated with the selected subgraph in hierarchical order. 4. Select an applicable graph. 5. Select <i>View details</i>.
Filter graph instances in the <i>Instances</i> tab.	<p>Use the filter tool to control the graphs to view:</p> <ol style="list-style-type: none"> 1. Open the <i>Instances</i> tab. 2. Select the Filter <i>Add Filter</i> icon. 3. Define the required filter conditions. Filter the listed graphs based on attributes, such as the source name, execution status, time of execution and more. 4. To filter and view graph instances executed on the same day, hour, or week, select <i>Hour</i>, <i>Day</i>, or <i>Week</i> in the menu bar and the view changes to records for that time frame. <p>If you're a tenant administrator, you can modify the filter to view graph instances of other users.</p> <p>Developer member users can modify the filter to view graph instances of other users when they are assigned the <code>sap.dh.monitoring</code> policy.</p>

Actions	Description
Open source graph in the <i>Instances</i> tab.	<p>For any selected graph instance, launch the source graph in the Modeler application from the Monitoring application.</p> <ol style="list-style-type: none"> 1. Open the <i>Instances</i> tab. 2. Select an instance. The application opens a graph overview pane at right that displays the source graph of the selected graph instance. 3. Select <i>Open Source Graph</i> in the menu bar of the overview pane. The Monitoring application opens the Modeler application in a new browser tab with the source graph of the selected instance opened in the graph editor. 4. View or edit the graph configuration.
View graph configurations in the <i>Instances</i> tab.	<p>For any selected graph instance, you can view the configurations defined for its source graph.</p> <ol style="list-style-type: none"> 1. Open the <i>Instances</i> tab. 2. Select an instance. The application opens a graph overview pane at right that displays the source graph of the selected graph instance. 3. Select the  <i>Show Configuration</i> icon. The application opens the <i>Configuration</i> pane at right where you can view read-only graph configuration information. 4. Optional: Right-click an operator in the graph overview pane and choose <i>Open Configuration</i>. The <i>Configuration</i> pane shows read-only operator configuration information.
Stop a graph instance execution in the <i>Instances</i> tab.	<p>Stop the execution of a graph instance when it's the running or pending state.</p> <ol style="list-style-type: none"> 1. Open the <i>Instances</i> tab. 2. Select a graph instance that has a running status. The graph overview pane opens at right. 3. Select the  <i>Stop Execution</i> icon in the menu bar.
Search graph instances in the <i>Instances</i> tab.	<p>To search for any graph instance, use the search bar in the <i>Instances</i> tab. Search for any graph instance based on the instance name or its source graph name.</p>
Archive graph instance in the <i>Instances</i> tab.	<p>Archive a completed or dead graph instance from the Pipeline Engine only.</p> <ol style="list-style-type: none"> 1. Open the <i>Instances</i> tab. 2. Select the  <i>Select an Action</i> icon in the <i>Actions</i> column of the applicable instance. 3. Select <i>Archive</i>.

Note

The archived instances remain in the system for a default period of 90 days. The tenant administrator can configure the retention time via the System Management application.

Actions	Description
Download diagnostics information and view logs in the Instances tab.	<p>Use the Monitoring application to download graph diagnostics information and view logs to help diagnose and solve errors with graphs.</p> <p>View logs generated for certain operators, such as the data workflow operator, when you execute the graph.</p> <ol style="list-style-type: none"> 1. Open the Instances tab. 2. Select the ... Select an Action icon in the Actions column of the applicable instance. 3. Select Download Diagnostic Info. <p>The application downloads a zipped archive of graph information automatically.</p> <p>For certain operators, such as the data workflow operators, view logs generated for the operator execution. To view these logs:</p> <ol style="list-style-type: none"> 1. Open the Instances tab. 2. Select the graph. The graph overview pane opens at right. 3. Select the Process Logs tab in the lower pane. The Logs tab contains a list of logs for the selected graph. 4. Use the filter lists at the top of the tab to filter for specific groups and processes, or search for the specific log. Scroll through the processes to read the log texts.
Manage schedules in the Schedules tab.	<p>View graphs that are scheduled for execution and create, edit, or delete schedules.</p> <p>For more information about scheduling graphs for execution, see “Schedule Graph Executions” in the <i>Modeling Guide</i>.</p> <p>A tenant administrator can view their schedules and the schedules of other users in the Monitoring application. If you're a tenant administrator, you can perform the following tasks in the Schedules tab on other users' schedules:</p> <ul style="list-style-type: none"> • View schedules by applying filters. • Edit or stop schedules. • Suspend or resume schedules. • View executed instances of schedules.

Actions	Description
Monitor the execution of replication flows in the Replications tab.	View information about replication data flows, including connection details, load progress, and user information.









For tasks, the [Replications](#) tab displays information, such as the number of tasks in different states, priority settings, number of operations and partitions, and execution timestamps. (The number of operations displayed may be higher than the actual amount of records transferred as data may be transferred twice in case of error situations.) Use the [Search](#) box in the menu bar to filter the list of replications on any metadata value.

Note

When a source system is unreachable due to network issues or down due to maintenance, replication tasks executed at that time might fail after exhausting retries. If this happens, you must manually resume the replication tasks.

To avoid this situation, we recommend that you suspend any running Replication flows before doing maintenance and then resume them once maintenance is done.

Select a replication flow to perform the following tasks:

- Display the replication in the Modeler application by selecting [Go to Monitoring](#) in the menu bar.
- Select the  [Settings](#) icon to configure the following settings:
 - [Replication Priority](#): Select [High](#), [Medium](#), or [Low](#). Selecting [High](#) runs this replication flow before others. The default is [Medium](#).
 - [Source Maximum Connections](#): Increase or limit the number of source connections permitted to the source. The default is 10.
 - [Target Maximum Connections](#): Increase or limit the number of target connections permitted to the target. The default is 10.
- To start the replication flow, select the  [Start Execution](#) icon.
- To suspend the execution of the replication flow, select the  [Suspend Execution](#) icon.
- To refresh the list of replication flows, select the  [Refresh](#) icon.
- Select a replication to display the tasks in the lower pane. Perform the following tasks in the [Tasks](#) area, after selecting a task:
 - Select the  [Settings](#) icon to configure [Task Priority](#) for the task. For example, select [High](#) to run this task before others. The default is [Medium](#).
 - Select the  [Start Execution](#) icon to run the task.
 - Select the  [Suspend Execution](#) icon to suspend execution of the task.
 - Select the  [Refresh](#) icon to refresh the list of tasks.

Related Information

[Schedule Graph Executions](#)

5.4 Stop, Restart, and Pause Tenant User Graphs

Manage your tenant workload by using the *Stop*, *Restart*, and *Pause* options for your tenant user graphs (pipelines).

Prerequisites

Only tenant administrators with the `sap.dh.admin` policy can manage the tenant workload by controlling tenant user graphs executions. To free resources for a higher-priority task, use the *Stop* or *Pause* options to temporarily stop or pause a running graph. After the important task completes, use the *Restart* option on the stopped or paused graph. You can also use the *Restart* option on dead graphs.

Context

To manage the tenant workload, perform the following steps:

Procedure

1. Log into SAP Data Intelligence and select the *Data Intelligence Monitoring* tile.

The Monitoring application opens.

2. Open the *Instances* tab.
3. Stop a running graph:
 - a. Choose the applicable running graph under the *Name* column.

The selected graph opens.

- b. Select  (*Stop*) from the menu bar.

4. Pause a running graph:
 - a. Choose the applicable running graph under the *Name* column.

The selected graph opens.

- b. Select  (*Pause*) in the menu bar.

5. Restart a paused, stopped, or dead graph:
 - a. Choose the paused, stopped, or dead graph from the *Name* column.

The selected graph opens.

- b. Select *Navigate to Latest Instance* in the menu bar.

Note

If the graph has been restarted previously, the option *Navigate to the Latest Instance* is active. If the option *Navigate to Latest Instance* isn't active, the opened graph is the latest instance.

- c. Select  (*Restart*) in the menu bar.

5.5 Accessing the SAP Data Intelligence Monitoring Query API

Tenant administrators can access the SAP Data Intelligence Monitoring Query API to retrieve tenant application performance metrics from the Prometheus monitoring service running in an SAP Data Intelligence cluster.

The SAP Data Intelligence Monitoring Query API complements the SAP Data Intelligence Monitoring application. It provides access for external monitoring and visualization tools to a collection of performance metrics of the Kubernetes pods running in a tenant:

- Pod CPU usage
- Pod memory usage
- Pod network usage
- Pod readiness status

The SAP Data Intelligence Monitoring Query API implements the HTTPS endpoint `/app/diagnostics-gateway/monitoring/query/api/v1`. For example, if the address of the SAP Data Intelligence cluster running your tenant is `dataintelligence.example.com`, then the SAP Data Intelligence Monitoring Query API endpoint is available at `https://dataintelligence.example.com/app/diagnostics-gateway/monitoring/query/api/v1`.

The endpoint implements the paths and HTTP methods in the following table:

Paths	Description	Methods	Required Parameters	Optional Parameters
<code>/query</code>	PromQL instant query	GET, POST	query	time
<code>/query_range</code>	PromQL range query	GET, POST	query, start, end, step	
<code>/series</code>	PromQL series request	GET, POST	match[]	start, end

This documentation describes the recommended usage and integration patterns for the SAP Data Intelligence Monitoring Query API, focusing on the supported query expressions for the parameters `query` and `match[]`. For a detailed specification of the SAP Data Intelligence Monitoring Query API, see the [SAP Business Accelerator Hub](#).

The query expressions supported for the parameters `query` and `match[]` are a subset of the [Prometheus Version 2.x Query Language](#) (PromQL). In addition to being proper PromQL expressions, supported query expressions must conform with the following restrictions:

- Only expressions considered stable according to the [Prometheus Version 2.x API Stability Guarantees](#) are supported. This explicitly excludes the `holt_winters` PromQL function.
- Only expressions that satisfy the following tenant access restrictions are supported.

Note

The SAP Data Intelligence Monitoring Query API is intended to behave similar to the respective endpoints of the [Prometheus Version 2.x HTTP API](#). Functionality of the Prometheus Version 2.x HTTP API that is

accessible but not documented in this guide is not part of the SAP Data Intelligence Monitoring Query API and may be restricted or removed in future versions or deployments of SAP Data Intelligence.

The following table lists the currently available pod performance metrics:

Pod Performance Metrics Names

Metric Name	Description
sap_pod_cpu_usage_seconds_total	Cumulative pod CPU core usage in seconds since pod start
sap_pod_memory_working_set_bytes	Current pod memory usage in bytes
sap_pod_network_bytes_total	Cumulative pod network usage in bytes since pod start
sap_pod_status_ready	Pod readiness status, 1: ready, 0: not ready

The following table lists the labels of the currently available pod performance metrics:

Pod Performance Metrics Labels

Metric Label	Description	Example
__name__	Metric name	sap_pod_memory_working_set_bytes
access_category	Metric access category	pod-performance
datahub_sap_com_app	Pod cluster application name	vflow
datahub_sap_com_app_component	Pod cluster component name	execution
graph	Pipeline graph ID	40f633c143954ad78e5460598fa55d43
vflow_datahub_sap_com_graph_source	Pipeline graph source	com-sap-demo-datagenerator
vsystem_datahub_sap_com_component	Pod system management component name	workload
vsystem_datahub_sap_com_tenant	Tenant name	default
vsystem_datahub_sap_com_tenant_uid	Tenant UID	744e5b4f8f4e49579038c820933c3f28
vsystem_datahub_sap_com_user	Username	member
namespace	Kubernetes namespace	dataintelligence
pod_name	Pod name	vflow-graph-40f633c143954ad78e5460598fa55d43-d-php9k2b4t2vhdtmx

For a detailed explanation and common usage patterns of these metrics and labels, see [Advanced Query Expressions \[page 248\]](#).

Note

The tenant metrics provided by the SAP Data Intelligence Monitoring Query API are considered content and are not part of the API. They depend on the specific SAP Data Intelligence instance and the running applications. Upon updates of SAP Data Intelligence, metrics may change in scope and specification (metric names and labels). For information about retrieving the metrics currently available for your tenant, see [Using a Series Request to List Available Metrics \[page 246\]](#).

Related Information

[Access Restrictions \[page 240\]](#)

[Metric Resolution and Retention \[page 240\]](#)

[Retrieving Your Tenant UID \[page 241\]](#)

[Testing the SAP Data Intelligence Monitoring Query API \[page 242\]](#)

[Running PromQL Instant Queries \[page 243\]](#)

[Running PromQL Range Queries \[page 244\]](#)

[Using a Series Request to List Available Metrics \[page 246\]](#)

[Accessing the SAP Data Intelligence Monitoring Query API Via POST Requests \[page 247\]](#)

[Advanced Query Expressions \[page 248\]](#)

[Advanced Query Expressions \[page 248\]](#)

[Using a Series Request to List Available Metrics \[page 246\]](#)

[SAP Data Intelligence Cloud API !\[\]\(a946603c62fef75051895d7de0479d08_img.jpg\)](#)

[Prometheus Version 2.x Query Language !\[\]\(33c10885660043ca3b8430b563a5ce47_img.jpg\)](#)

[Prometheus Version 2.x API Stability Guarantees !\[\]\(4b9475ed8ef390cd64e4284ba68752f7_img.jpg\)](#)

[Prometheus Version 2.x HTTP API !\[\]\(b6e0e7cdae4a4db39af63849547dcd54_img.jpg\)](#)

5.5.1 Access Restrictions

Access to the SAP Data Intelligence Monitoring Query API is restricted to tenant administrators; that is, tenant users that have been assigned the `sap.dh.admin` policy.

In addition, access to SAP Data Intelligence metrics is restricted by access category and tenant. Tenant users can only access metrics for their own tenant. Any metric included in a PromQL expression must specify the following two labels:

- `access_category`: "pod-performance"
- `vsystem_datahub_sap_com_tenant_uid`: <tenant UID>

For more information about retrieving the tenant UID of your tenant, see [Retrieving Your Tenant UID \[page 241\]](#).

5.5.2 Metric Resolution and Retention

SAP Data Intelligence application metrics are intended for real time usage in external monitoring, visualization, and alerting systems.

These metrics are deleted by volume limits and a pre-set retention period. Retention of application metrics is performed by deleting the oldest samples of each time series. The retention period is part of the infrastructure settings and is not configurable. It may be adapted by SAP at any time, but will always be at least 24 hours. If you require access to metrics older than 24 hours, use the SAP Data Intelligence Monitoring Query API to fetch the metrics for external persistence.

Due to the concepts and implementation of the underlying Prometheus service, SAP Data Intelligence application metrics have a fixed time-series resolution (sample frequency). The time-series resolution is part of the infrastructure settings and not configurable. It is typically set to one minute but may be increased up to ten minutes depending on the system load. Between two sample points, the instant value of a particular metric does not change. Consequently, querying the SAP Data Intelligence Monitoring Query API more than once per minute for the same value is not supported. By the nature of the Prometheus metric collection mechanism, this also means that metric values that occur between the most recent regular sample point and the shutdown time of a pod are not recorded.

5.5.3 Retrieving Your Tenant UID

Accessing the metrics of your tenant via the SAP Data Intelligence Monitoring Query API requires knowledge of your tenant UID.

Context

The tenant UID is different from the tenant name. You can obtain it using the `vctl` tool with the following command:

```
# login to SAP Data Intelligence as tenant admin
vctl login [...]
# get the detailed tenant information, including the tenant ID
vctl tenant get [EDIT HERE: tenant name] -o yaml
```

The result should look similar to the following output:

```
apiVersion: v3
[...]
id: 744e5b4f8f4e49579038c820833c3f28
name: default
status: Ready
[...]
```

In the example, the `name` field value `default` is the tenant name used to authenticate against SAP Data Intelligence, and the `ID` field `744e5b4f8f4e49579038c820833c3f28` is the tenant UID required in the PromQL expression when accessing the SAP Data Intelligence Monitoring Query API.

5.5.4 Testing the SAP Data Intelligence Monitoring Query API

To test the SAP Data Intelligence Monitoring Query API, send a simple constant term expression for which the result is known.

Procedure

1. As a prerequisite, expose the SAP Data Intelligence cluster address, the UID of your tenant, the tenant name, and your user name as shell variables:

```
SAP_DI_CLUSTER_ADDRESS="[EDIT HERE: SAP Data Intelligence cluster address]"
SAP_DI_TENANT_UID="[EDIT HERE: tenant UID]"
SAP_DI_TENANT="[EDIT HERE: tenant name]"
SAP_DI_USER="[EDIT HERE: username]"
```

2. Send a request to test the SAP Data Intelligence Monitoring Query API using the constant term PromQL expression `1+1`. This request uses basic authentication with `<tenant>\<user>` as user, and you must enter your password when prompted.

```
curl -G -u "${SAP_DI_TENANT}\/\${SAP_DI_USER}" "https://${SAP_DI_CLUSTER_ADDRESS}/app/diagnostics-gateway/monitoring/query/api/v1/query" --data-urlencode "query=1+1"
```

Note

You can also execute the query using the user interface of the [SAP Business Accelerator Hub](#).

After you enter the password, the result should look similar to the following output (formatted with line breaks and indents for improved readability):

```
{
  "status": "success",
  "data": {
    "resultType": "scalar",
    "result": [ 1609498800.000, "2" ]
  }
}
```

Results

The status (`success` or `error`) indicates the execution status. The result of the constant term query is the scalar value `2`, which is indeed the result of the expression `1+1`. The number `1609498800.000` in the query result is the Unix timestamp of the current time (in the example, corresponding to `2021-01-01T12:00:00.000z` in RFC 3339).

5.5.5 Running PromQL Instant Queries

A PromQL instant query returns the metric value of a set of metrics for a given time (by default for the current time).

Procedure

1. As a prerequisite, expose the SAP Data Intelligence cluster address, the UID of your tenant, the tenant name, and your user name as shell variables:

```
SAP_DI_CLUSTER_ADDRESS="[EDIT HERE: SAP Data Intelligence cluster address]"
SAP_DI_TENANT_UID="[EDIT HERE: tenant UID]"
SAP_DI_TENANT="[EDIT HERE: tenant name]"
SAP_DI_USER="[EDIT HERE: username]"
```

2. Send an instant query request to query a pod performance metric (in the following example, the pod memory usage in bytes) of your tenant:

```
curl -G -u "$
{SAP_DI_TENANT}\\$SAP_DI_USER" "https://$SAP_DI_CLUSTER_ADDRESS/app/
diagnostics-gateway/monitoring/query/api/v1/query" --data-
urlencode "query=sap_pod_memory_working_set_bytes{access_category=\"pod-
performance\",vsystem_datahub_sap_com_tenant_uid=\"$SAP_DI_TENANT_UID}\""
```

Note

You can also execute the query using the user interface of the [SAP Business Accelerator Hub](#).

The result should look similar to the following output (formatted with line breaks and indents for improved readability):

```
{
  "status": "success",
  "data": {
    "resultType": "vector",
    "result": [
      {
        "metric": {
          "__name__": "sap_pod_memory_working_set_bytes",
          "access_category": "pod-performance",
          "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
          "pod_name": "vflow-graph-40f633c143954ad78e5460598fa55d43-d-
php9k2b4t2vhdtmx"
          [...]
        },
        "value": [ 1609498800.000, "481116160" ]
      },
      [...]
    ]
  }
  "metric": {
    "__name__": "sap_pod_memory_working_set_bytes",
    "access_category": "pod-performance",
    "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
    "pod_name": "vflow-graph-fcfe1d5f4bde4b25bec7dcded118dc82-c-
lg4plpzvplxfcmv5"
```

```

    [ ... ]
    },
    "value": [ 1609498800.000, "101072896" ]
  ]
}
}
}

```

Results

The result represents an instant vector. Each element of the vector has a `metric` field that contains all of the labels that describe the metric instant (including the special metric name label `__name__`). For each metric, there is a dedicated sample (UNIX timestamp and actual metric value). The metric values are produced for the current time. To query the metric instant values at another point in time, specify the optional HTTP parameter `time` in RFC 3339 format.

For details about instant query expressions, see [Prometheus Version 2.x Query Language](#).

5.5.6 Running PromQL Range Queries

A PromQL range query returns a series of metric samples for a specified time range.

Procedure

1. As a prerequisite, expose the SAP Data Intelligence cluster address, the tenant UID of your tenant, the tenant name, and your user name as shell variables:

```

SAP_DI_CLUSTER_ADDRESS="[EDIT HERE: SAP Data Intelligence cluster address]"
SAP_DI_TENANT_UID="[EDIT HERE: tenant UID]"
SAP_DI_TENANT="[EDIT HERE: tenant name]"
SAP_DI_USER="[EDIT HERE: username]"

```

2. Send a range query request to query a pod performance metric (in the following example, the pod memory usage in bytes) of your tenant. You must specify the start and end times of the query range (in RFC 3339 or Unix Timestamp format) and a step size in the PromQL duration format. The duration format is a string consisting of an integer and a unit with no spaces between them.

The supported units are:

- `ms`: milliseconds
- `s`: seconds
- `m`: minutes
- `h`: hours
- `d`: days (equals 24h)
- `w`: weeks (equals 7d)

- y: years (equals 365d)

The minimum supported step size is one minute.

```

QUERY_START="[EDIT HERE: replace date in 2021-01-01T12:00:00.000Z by
yesterday]"
QUERY_END="[EDIT HERE: replace date in 2021-01-01T12:00:00.000Z by tomorrow]"
QUERY_STEP="30m"
curl -G -u "$
${SAP_DI_TENANT}\\${SAP_DI_USER}" "https://${SAP_DI_CLUSTER_ADDRESS}/app/
diagnostics-gateway/monitoring/query/api/v1/query_range" --data-
urlencode "query=sap_pod_memory_working_set_bytes{access_category=\`pod-
performance\`,vsystem_datahub_sap_com_tenant_uid=\`${SAP_DI_TENANT_UID}\`}" --
data-urlencode "start=${QUERY_START}" --data-urlencode "end=${QUERY_END}" --
data-urlencode "step=${QUERY_STEP}"

```

Note

You can also execute the query using the user interface of the [SAP Business Accelerator Hub](#).

The result should look similar to the following output (formatted with line breaks and indents for improved readability):

```

{
  "status": "success",
  "data": {
    "resultType": "matrix",
    "result": [
      {
        "metric": {
          "__name__": "sap_pod_memory_working_set_bytes",
          "access_category": "pod-performance",
          "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
          "pod_name": "vflow-graph-40f633c143954ad78e5460598fa55d43-d-
php9k2b4t2vhdtmx"
        },
        "values": [
          [ 1609498800.000, "481116160" ],
          [ 1609498830.000, "481116283" ],
          [ 1609498860.000, "481116182" ]
        ]
      },
      [...]
      {
        "metric": {
          "__name__": "sap_pod_memory_working_set_bytes",
          "access_category": "pod-performance",
          "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
          "pod_name": "vflow-graph-fcfeld5f4bde4b25bec7dcded118dc82-c-
lg4plpzzvplxfcmv5"
        },
        "values": [
          [ 1609498800.000, "101072896" ],
          [ 1609498830.000, "101072896" ],
          [ 1609498860.000, "101072896" ]
        ]
      }
    ]
  }
}

```

Results

The result represents a range vector. Like an instant vector, each element of the range vector has a `metric` field that contains all of the labels that describe the metric. For each metric, there is a dedicated series of samples in line with the start time, end time, and step size.

For details on range query expressions, see [Prometheus Version 2.x Query Language](#).

5.5.7 Using a Series Request to List Available Metrics

A PromQL series request returns a list of available metrics.

Procedure

1. As a prerequisite, expose the SAP Data Intelligence cluster address, the UID of your tenant, the tenant name, and your user name as shell variables:

```
SAP_DI_CLUSTER_ADDRESS="[EDIT HERE: SAP Data Intelligence cluster address]"
SAP_DI_TENANT_UID="[EDIT HERE: tenant UID]"
SAP_DI_TENANT="[EDIT HERE: tenant name]"
SAP_DI_USER="[EDIT HERE: username]"
```

2. Send the following series request to list all metrics that are currently available.

Note

The parameter that specifies the PromQL expression is `match[]`, not `query`.

```
curl -G -u "$
{SAP_DI_TENANT}\\${SAP_DI_USER}" "https://${SAP_DI_CLUSTER_ADDRESS}/app/
diagnostics-gateway/monitoring/query/api/v1/series"
--data-urlencode "match[ ]={access_category=\"pod-
performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}"
```

Note

You can also execute the query using the user interface of the [SAP Business Accelerator Hub](#).

The result should look similar to the following output (formatted with line breaks and indents for improved readability):

```
{
  "status": "success",
  "data": [
    {
      "__name__": "sap_pod_cpu_usage_seconds_total",
      "access_category": "pod-performance",
      "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
      "pod_name": "vflow-graph-40f633c143954ad78e5460598fa55d43-d-
php9k2b4t2vhdtmx"
```

```

    [...]
  },
  {
    "__name__": "sap_pod_memory_working_set_bytes",
    "access_category": "pod-performance",
    "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
    "pod_name": "vflow-graph-40f633c143954ad78e5460598fa55d43-d-
php9k2b4t2vhdtmx"
    [...]
  },
  {
    "__name__": "sap_pod_network_bytes_total",
    "access_category": "pod-performance",
    "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
    "pod_name": "vflow-graph-40f633c143954ad78e5460598fa55d43-d-
php9k2b4t2vhdtmx"
    [...]
  },
  {
    "__name__": "sap_pod_status_ready",
    "access_category": "pod-performance",
    "vsystem_datahub_sap_com_tenant_uid":
"744e5b4f8f4e49579038c820933c3f28",
    "pod_name": "vflow-graph-40f633c143954ad78e5460598fa55d43-d-
php9k2b4t2vhdtmx"
    [...]
  },
  [...]
]
}

```

Results

The result lists all of the available metrics represented by their labels. The special metric name label `__name__` itemizes all available metrics, which can be accessed using instant or range queries:

- `sap_pod_memory_working_set_bytes`
- `sap_pod_cpu_usage_seconds_total`
- `sap_pod_network_bytes_total`
- `sap_pod_status_ready`

Optional query parameters `start` and `end` can be specified to list all metrics that were available during an earlier period of time.

5.5.8 Accessing the SAP Data Intelligence Monitoring Query API Via POST Requests

You can send HTTP GET requests as `x-www-form-urlencoded` HTTP POST requests.

PromQL expressions sent to the SAP Data Intelligence Monitoring Query API can be complex and too extensive to be URL query-encoded in an HTTP GET request. Therefore, all HTTP GET requests described in [Running PromQL Instant Queries \[page 243\]](#), [Running PromQL Range Queries \[page 244\]](#), and [Using a Series Request](#)

to [List Available Metrics \[page 246\]](#) can also be sent as `x-www-form-urlencoded` HTTP POST requests. The requests have the same (read-only) semantics; the only difference is that the parameters are encoded in the request body rather than in the URL.

Note

SAP Data Intelligence implements the Cross-Origin Resource Sharing (CORS) concept to protect against Cross Site Request Forgery (CSRF) attacks. All HTTP POST requests to the SAP Data Intelligence Monitoring Query API must specify the header `x-Requested-With: fetch`. If this header is not specified, CORS-related error messages, such as the following `curl` error message, will occur:

```
Forbidden cross-site request
```

To replace an HTTP GET request to the SAP Data Intelligence Monitoring Query API by an equivalent HTTP POST request, replace the parameter `-g` with the parameters `-X POST -H "x-requested-with: fetch"` in the `curl` command of the documented example.

5.5.9 Advanced Query Expressions

Here we discuss advanced PromQL queries to the SAP Data Intelligence Monitoring Query API that you can use to monitor your SAP Data Intelligence tenant.

To execute the queries, execute the following command after substituting the respective variables:

```
SAP_DI_CLUSTER_ADDRESS="[EDIT HERE: SAP Data Intelligence cluster address]"
SAP_DI_TENANT_UID="[EDIT HERE: tenant UID]"
SAP_DI_TENANT="[EDIT HERE: tenant name]"
SAP_DI_USER="[EDIT HERE: username]"
SAP_DI_QUERY="[EDIT HERE: select a query from the list of advanced query
expressions below]"
curl -X POST -H "x-requested-with: fetch" -u "${SAP_DI_TENANT}
\\${SAP_DI_USER}" "https://${SAP_DI_CLUSTER_ADDRESS}/app/diagnostics-gateway/
monitoring/query/api/v1/query" --data-urlencode "query=${SAP_DI_QUERY}"
```

Note

You can also execute the query using the user interface of the [SAP Business Accelerator Hub](#).

The queries presented in this section are based on the following pod performance metrics:

- `sap_pod_memory_working_set_bytes`
- `sap_pod_cpu_usage_seconds_total`
- `sap_pod_network_bytes_total`
- `sap_pod_status_ready`

Related Information

[Basic Pod Performance Metrics Usage \[page 249\]](#)

[Tenant Pod Performance \[page 250\]](#)

[User Pod Performance \[page 250\]](#)

[Pipeline Graph Performance \[page 251\]](#)

5.5.9.1 Basic Pod Performance Metrics Usage

Advanced PromQL queries to the SAP Data Intelligence Monitoring Query API that monitor basic pod performance metrics usage.

Pod memory usage in bytes

```
SAP_DI_QUERY="sap_pod_memory_working_set_bytes{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}"
```

Pod CPU cores usage

```
SAP_DI_QUERY="rate(sap_pod_cpu_usage_seconds_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}[5m])"
```

The `rate` function is applied because the base metric `sap_pod_cpu_usage_seconds_total` records the total CPU usage over the lifetime of a pod and is not very informative on its own. For the expression above, a value of 0.1 corresponds to an average usage of 1/10th of the CPU time of a single core over the past five minutes, while a value of 2 corresponds to an average usage of two full CPU cores.

Pod network usage as bytes per second

```
SAP_DI_QUERY="rate(sap_pod_network_bytes_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}[5m])"
```

The `rate` function is applied because the base metric `sap_pod_network_bytes_total` records the total network usage over the lifetime of a pod and is not very informative on its own. The network usage as bytes per second is computed over the past five minutes.

Pod readiness status

```
SAP_DI_QUERY="sap_pod_status_ready{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}"
```

This is a zero-one metric. The value 1 represents a ready pod; the value 0 represents a non-ready pod.

Smoothed pod readiness status

```
SAP_DI_QUERY="avg_over_time(sap_pod_status_ready{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}[5m])"
```

The result is a sliding window average of the pod readiness status over the past five minutes (based on four or five samples due to a sample resolution of one minute). The values lie between zero (pod not ready for the past five minutes) and one (pod ready for the past five minutes). This metric is suitable to define an alert threshold, for example, at 0.7, allowing the pod to be not ready for a minute during restarts without raising an alert.

Note

Do not set the time interval for `rate` or `avg_over_time` functions in your PromQL expressions below 5m. The minimum time series resolution of the queried metrics is at least one minute (see [Metric Resolution and Retention \[page 240\]](#)). This means that for an interval of five minutes, the `rate` or `avg_over_time`

functions are already based only on four samples points. Reducing the interval further may result in too few samples to calculate these functions.

5.5.9.2 Tenant Pod Performance

Advanced PromQL queries to the SAP Data Intelligence Monitoring Query API that monitor tenant pod performance.

Total memory usage in bytes of all pods of the tenant

```
SAP_DI_QUERY="sum(sap_pod_memory_working_set_bytes{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"})"
```

Total CPU cores usage of all pods of the tenant

```
SAP_DI_QUERY="sum(rate(sap_pod_cpu_usage_seconds_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}[5m]))"
```

Total network usage as bytes per second of all pods of the tenant

```
SAP_DI_QUERY="sum(rate(sap_pod_network_bytes_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"}[5m]))"
```

Total pod count

```
SAP_DI_QUERY="count(sap_pod_status_ready{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"})"
```

Total count of ready pods

```
SAP_DI_QUERY="sum(sap_pod_status_ready{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\"})"
```

5.5.9.3 User Pod Performance

Advanced PromQL queries to the SAP Data Intelligence Monitoring Query API that monitor user pod performance.

Total memory usage in bytes for each user

```
SAP_DI_QUERY="sum(sap_pod_memory_working_set_bytes{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",vsystem_datahub_sap_com_user!=\"\"}) by (vsystem_datahub_sap_com_user)"
```

Total CPU cores usage for each user

```
SAP_DI_QUERY="sum(rate(sap_pod_cpu_usage_seconds_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",vsystem_datahub_sap_com_user!=\"\"}[5m])) by (vsystem_datahub_sap_com_user)"
```

Total network usage as bytes per second for each user

```
SAP_DI_QUERY="sum(rate(sap_pod_network_bytes_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",vsystem_datahub_sap_com_user!=\"\"}[5m])) by (vsystem_datahub_sap_com_user)"
```

Total pod count for each user

```
SAP_DI_QUERY="count(sap_pod_status_ready{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",vsystem_datahub_sap_com_user!=\"\"}) by (vsystem_datahub_sap_com_user)"
```

5.5.9.4 Pipeline Graph Performance

Advanced PromQL queries to the SAP Data Intelligence Monitoring Query API that monitor pipeline graph performance.

Total memory usage in bytes of all pods for each (multi-pod) graph

```
SAP_DI_QUERY="sum(sap_pod_memory_working_set_bytes{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",graph!=\"\"}) by (graph)"
```

Total CPU cores usage of all pods for each (multi-pod) graph

```
SAP_DI_QUERY="sum(rate(sap_pod_cpu_usage_seconds_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",graph!=\"\"}[5m])) by (graph)"
```

Total network usage as bytes per second for each (multi-pod) graph

```
SAP_DI_QUERY="sum(rate(sap_pod_network_bytes_total{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",graph!=\"\"}[5m])) by (graph)"
```

Readiness status for each (multi-pod) graph

```
SAP_DI_QUERY="max(sap_pod_status_ready{access_category=\"pod-performance\",vsystem_datahub_sap_com_tenant_uid=\"${SAP_DI_TENANT_UID}\",graph!=\"\"}) by (graph)"
```

6 Integration Monitoring in SAP Cloud Application Lifecycle Management

Register your SAP Data Intelligence tenant with SAP Cloud Application Lifecycle Management (SAP Cloud ALM) and use Integration Monitoring to monitor your data exchange processes.

Prerequisites

Note

Integration Monitoring is not supported when using Cluster Hibernation and Wakeup.

To configure SAP Cloud ALM Integration Monitoring for your SAP Data Intelligence tenant, you will require the following information:

- Set up SAP Cloud ALM: Request an SAP Cloud ALM provisioning as described in [Requesting SAP Cloud ALM](#). An SAP Cloud ALM tenant will be created for you.

Note

If you already have access to SAP Cloud ALM and the information above, you do not need to create a new SAP Cloud ALM tenant.

- Download the SAP Cloud ALM service key. To download the SAP Cloud ALM service key, follow the steps described in [Retrieving SAP Cloud ALM Service Key](#).
- The service key of this tenant includes the following information required later in the process:
 - CALM_API_ENDPOINT: application base URL (endpoints.Api)
 - CALM_CLIENT_ID: clientid
 - CALM_CLIENT_SECRET: clientsecret
 - CALM_AUTHENTICATION_ENDPOINT: oAuth URL

Context

Integration Monitoring provides end-to-end visibility in interface calls and message flows, including possible technical root causes across cloud services and systems involved with SAP Data Intelligence.

Integration Monitoring for SAP Data Intelligence provides an overview of the incoming and outgoing messages that have SAP Passport information of your SAP Data Intelligence cluster.

For more information, see the SAP Cloud ALM Expert Portal under [SAP Cloud ALM](#) and [Integration & Exception Monitoring](#).

Procedure

1. Log into the SAP Data Intelligence System Management of your tenant. You can do this using either the SAP Data Intelligence System Management user interface (see [Access SAP Data Intelligence Cloud System Management \[page 27\]](#)) or the System Management Command-Line Client for SAP Data Intelligence (see [System Management Command-Line Client Reference for SAP Data Intelligence](#)).
2. Create the technical user `sap-cloud-alm-technical-user`. You can do this using either the SAP Data Intelligence System Management user interface (see [Create and Delete users and User Instances \[page 47\]](#)) or the `vctl` command-line client:

```
vctl user create <TENANT> sap-cloud-alm-technical-user <password> member
```

3. Create a new policy for connections and assign it to the technical user.

The policy is called `connectionConfigurationAllRW`. Add two resources as follows respectively:

- Type `connectionConfiguration`
- Activities `read` and `write`

Set `ConnectionId` to `*`. Assign the policy to `sap-cloud-alm-technical-user`.

Alternatively, you can use the command-line client tool. Open an editor and create a `connectionConfigurationAllRWPolicy.json` file with the following content:

```
{
  "id": "connectionConfigurationAllRW",
  "description": "Read/Write access to all connection configurations in the
tenant",
  "enabled": true,
  "exposed": true,
  "resources": [
    {
      "resourceType": "connectionConfiguration",
      "contentData": { "activity": "read", "technicalName": "*" }
    },
    {
      "resourceType": "connectionConfiguration",
      "contentData": { "activity": "write", "technicalName": "*" }
    }
  ]
}
```

4. To assign the policy to the technical user, execute the following command:

```
vctl policy assign connectionConfigurationAllRW sap-cloud-alm-technical-user
```

5. Create an SAP Cloud ALM configuration with the technical user. You must provide the following information:
 - An endpoint to the SAP Cloud ALM API; for example, `https://eu10.alm.cloud.sap`
 - An authentication endpoint; for example, `https://calm-prod-eu10-rel-cloud-lob-integration.authentication.eu10.hana.ondemand.com/oauth/token`
 - A valid SAP Cloud ALM client ID and secret. You must create the configuration via a call to the SAP Data Intelligence API.
6. Open an editor and create the `calm-conf.json` file:

```
{
  "technicalName": "SAP_CLOUD_ALM",
  "typeId": "SAPCloudALM",
}
```

```

"typeVersion": "1.0.0",
"properties": {
  "apiEndpoint": <SAP Cloud ALM API endpoint>,
  "credentials": {
    "oauth2TokenEndpoint": <SAP Cloud ALM authentication endpoint>,
    "oauth2ClientId": <SAP Cloud ALM client id>,
    "oauth2ClientSecret": <SAP Cloud ALM client secret>
  }
}
}

```

7. Upload the `calm-conf.json` configuration file to the SAP Data Intelligence Connection Service by sending a POST request with the following data, using your preferred API client for REST queries.

Field	Value
HTTP method	POST
Hostname	Cluster address; for example, <code>https://vsystem.ingress...hana.ondemand.com</code>
Path	<code>/app/connection-service/v1/configurations</code>
User	<code><TENANT>\sap-cloud-alm-technical-user</code>
Header	Content-Type: application/json X-Requested-With: fetch
Body	<code><contents of calm-conf.json></code>

With curl, the command would look like this:

```

curl -X POST -H "x-requested-with: fetch" -H "Content-Type: application/json" -u "<TENANT>\\sap-cloud-alm-technical-user" "https://<SAP_DI_CLUSTER_ADDRESS>/app/connection-service/v1/configurations" --data-binary @calm-conf.json

```

Note

Depending on your shell, the exact syntax of the curl command may slightly differ. In particular, the escape sequence (the double backslash "\\") separating the tenant and user in the curl username may not be required. In this case, use a single backslash "\".

8. You will be prompted for the password that you previously assigned to the technical user `sap-cloud-alm-technical-user`.

As a result, this command should output a JSON object with the unique ID of the configuration; for example, `{ "id": "1a2b3c4d-5678-9abc-1234-0123456789ab" }`. You can use the ID to verify the configuration afterwards, with the following REST API query (note that for security reasons, the client ID is abbreviated and the client secret not shown at all in the output).

Field	Value
HTTP method	GET
Hostname	Cluster address; for example, <code>https://vsystem.ingress...hana.ondemand.com</code>

Field	Value
Path	/app/connection-service/v1/configurations/{id}
User	<TENANT>\sap-cloud-alm-technical-user
Header	Content-Type: application/json X-Requested-With: fetch
Body	<empty>

Again, using *curl*, the command should look like this (again, with a double backslash for unescaping):

```
curl -X GET -H "Content-Type: application/json" -u "<TENANT>\\sap-cloud-alm-technical-user" "https://<SAP_DI_CLUSTER_ADDRESS>/app/connection-service/v1/configurations/<id>"
```

Related Information

[Using the SAP Cloud ALM Integration Monitoring for SAP Data Intelligence \[page 255\]](#)

6.1 Using the SAP Cloud ALM Integration Monitoring for SAP Data Intelligence

After you have configured SAP Cloud ALM Integration Monitoring for SAP Data Intelligence, you can view the incoming and outgoing messages that have SAP Passport information in the SAP Cloud ALM user interface.

Context

Note

It may take up to 15 minutes after you create the SAP Cloud ALM configuration until messages are sent.

Procedure

1. To use the SAP Cloud ALM Integration Monitoring user interface for SAP Data Intelligence, follow the instructions in the [SAP Cloud ALM Documentation](#).
2. Go to the [SAP Cloud ALM Integration Monitoring](#) dashboard and click the "Select a Scope" filter icon.

3. Choose **Services/Systems** > **Services** and **Service Type**, press *Go* in the upper right, and then *Apply* in the bottom right.

Example

Create and run a pipeline in the SAP Data Intelligence Modeler that transfers data to an external system. The data transfer displays as outgoing messages in the SAP Cloud ALM user interface. Create a graph in Modeler that transfers data to an external system. This causes passports to be sent to SAP Cloud ALM. The monitoring messages should be visible in the SAP Cloud ALM user interface. Note that it may take up to 15 minutes after creating the SAP Cloud ALM configuration until messages are sent.

7 Maintaining SAP Data Intelligence

This section describes how to maintain SAP Data Intelligence.

Related Information

[On-Demand Certificate Renewal \[page 257\]](#)

7.1 On-Demand Certificate Renewal

On-demand certificate renewal provides solution to clusters with expired certificates, compromised private keys and security vulnerabilities in crypto-libraries.

This feature can be triggered using the low-level tooling (dhinstaller). Dhinstaller has a subcommand "certificates renew" which performs the renewal operation.

Currently, an SAP Data Intelligence instance contains three certificate trees:

- root certificate for client certificate authentication
- root certificate for internal communication
- root certificate for nats

The client certificate tree is used for external communication. The internal certificate and nats trees are used for internal communication inside SAP Data Intelligence.

Prerequisite: To rotate root certificates, you must be a cluster administrator.

Related Information

[Functional Cluster \[page 258\]](#)

[Non-Functional Cluster \[page 258\]](#)

7.1.1 Functional Cluster

To perform certificate renewal in a functioning cluster, run the following command in a shell window:

```
kubectl -n datahub-system exec -it datahub-operator-0 -- dhinstaller
certificates renew --namespace [installation namespace] --registry [registry
used in installation] --stack "" --renew-nats-ca --renew-internal-ca
```

If the renewal of nats ca certificate is not needed, you can omit `--renew-nats-ca`.

If the renewal of internal ca certificate is not needed, you can omit `--renew-internal-ca`.

At least one of `--renew-nats-ca` and `--renew-internal-ca` should be passed as an argument.

If a timeout occurs while switching the runlevel to Stopped, the cluster may be in a non-functional state; continue with the following steps.

7.1.2 Non-Functional Cluster

If SAP Data Intelligence System Management is not functional (for example, the certificates are expired), enable the `--disable-vsystème-hooks` flag and provide the user information. Also, validating webhook should be disabled for this operation.

To perform certificate renewal in a functioning cluster, run the following command in a shell window:

```
kubectl get validatingwebhookconfigurations.admissionregistration.k8s.io
validating-webhook-configuration -o yaml > /tmp/validating-webhook-
configuration.yaml
kubectl delete validatingwebhookconfigurations.admissionregistration.k8s.io
validating-webhook-configuration
kubectl -n datahub-system exec -it datahub-operator-0 -- dhinstaller
certificates renew --namespace [installation namespace] --registry [registry
used in installation] --stack "" --renew-nats-ca --renew-internal-ca --
username [installation username] -p [installation password] --system-password
[installation system password] --default-tenant-name [default tenant name] --
disable-vsystème-hooks true
sed 's/Unknown/None/g' /tmp/validating-webhook-configuration.yaml > /tmp/
validating-webhook-configuration2.yaml
kubectl apply -f /tmp/validating-webhook-configuration2.yaml
```

If the renewal of nats ca certificate is not needed, you can omit `--renew-nats-ca`.

If the renewal of internal ca certificate is not needed, you can omit `--renew-internal-ca`.

At least one of `--renew-nats-ca` and `--renew-internal-ca` should be passed as an argument.

At least one of ```--renew-nats-ca``` and ```--renew-internal-ca``` should be passed as an argument.

8 Exporting Customer Data

You can export customer data from different SAP Data Intelligence components to a target store for various intended purposes.

For example, you can export to move, extract, or copy the customer data from one system to another, or to process the data for different business use cases. The following components of SAP Data Intelligence can contain customer data from which you can export and retrieve.

- Application Database
- SAP Data Intelligence
- Files

→ Remember

Some parts of the stored data are user-specific. Therefore, due to access isolation, only respective users can access and export those files. Additionally, only tenant admins can export customer data and files from application databases and tenant layers respectively.

In the current version, you must manually trigger all exports. Additionally, you must also trigger the export in all existing tenants to retrieve a full export of customer data. This section describes the steps that you must execute to export the data from different components.

Component	Steps
Application Database	<p>You can export the SAP HANA customer data from your landscape using the SAP Data Intelligence Customer Data Export application.</p> <ol style="list-style-type: none"> 1. Log in to SAP Data Intelligence Launchpad with the tenant admin credentials. 2. Start the <i>Customer Data Export</i> application. 3. In the menu bar, choose <i>Export Data</i>. 4. In the <i>Configure Export</i> dialog: <ol style="list-style-type: none"> 1. Select a connection type. The connections of the selected connection type are listed in the <i>Export Location</i> dropdown. <div data-bbox="667 667 1396 779" style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px;"> <p>Note You can export only to S3 connection type.</p> </div> <ol style="list-style-type: none"> 2. Select an existing connection and enter the bucket details. You can manage the connections in Connection Management application. <div data-bbox="667 862 1396 1003" style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px;"> <p>Note For connections with endpoint URLs containing region code, the URL should follow the pattern: <code>s3.<region>.amazonaws.com</code>.</p> </div> 5. Click <i>Export</i>. The application displays a confirmation message on the status of the data export. 6. The exported target folder links are displayed in the <i>Customer Data Export</i> application home page. 7. To view the exported data, use <i>Browse Connections</i> link in the Metadata Explorer. For more information about Metadata Explorer, see Using the Metadata Explorer.
SAP Data Intelligence Files	<p>SAP Data Intelligence System Management application supports exporting SAP Data Intelligence files as .tgz files. You can download the content of all files by executing the following steps in the SAP Data Intelligence System Management application.</p> <ol style="list-style-type: none"> 1. Log on to every tenant with tenant admin and export all files in <i>Tenant Workspace</i> 2. Log on to every user of every tenant and download all files in <i>My Workspace</i>. <p>In the <i>System Management</i> application, choose the <i>Files</i> tab and switch to the <i>Split</i> view to identify files in the <i>Tenant Workspace</i> and <i>My Workspace</i> separately.</p>

Related Information

[Log in to SAP Data Intelligence Customer Data Export \[page 261\]](#)

8.1 Log in to SAP Data Intelligence Customer Data Export

You can access the SAP Data Intelligence Customer Data Export application from the SAP Data Intelligence Launchpad or directly launch the application with a stable URL.

Procedure

1. Open the SAP Data Intelligence Launchpad URL in a browser.
The welcome screen appears where you can enter the log in credentials.
2. Log in to the SAP Data Intelligence Launchpad by providing the following:
 - Your tenant name
 - Your username
 - Your password

For newly created instance the tenant name is "default". For user name and password, use the credentials used while creating a service.

The SAP Data Intelligence Launchpad opens and displays the initial home page. User details and the tenant ID are displayed on click of the profile icon in the upper-right area of the screen. The home page displays the applications available in the tenant based on the polices assigned to the user.

Note

If you enter an incorrect password five consecutive times within a minute, your account will be temporarily locked for ten seconds until your next attempt.

3. On the home page, choose *Customer Data Export* application.
The application UI opens displaying the initial screen.

9 Improving Performance

Describes ways that you can improve SAP Data Intelligence performance.

Related Information

[Improving CDC Graph Generator Operator Performance \[page 262\]](#)

9.1 Improving CDC Graph Generator Operator Performance

The performance of the Change Data Capture (CDC) Graph Generator operator depends on the initial table size and how rapidly change data is generated in the source system.

To optimize and process different scenarios, the following are some helpful tips for using the CDC Graph Generator operator.

Initial Load Volume

If the source table is large, it may be helpful to partition the source data using the *Partition Specification* parameter. After the partitioning is defined based on the resources available, you can set the *Max Number of Loaders* to the appropriate value (the default is 8).

Delta Change Rate

For a rapidly changing table, it is recommended that you set *Delta Graph Run Mode* to *polling interval* with a short polling interval so that the operator fetches data rapidly as well.

For slowly a changing table, you can either set *Delta Graph Run Mode* to *Manual* and then schedule the graph periodically (for example, N hours) or set *Delta Graph Run Mode* to *polling interval* with a high polling interval.

When you use a polling interval, the graph runs. In manual mode, the operator stops processing, and you can terminate the graph using a graph terminator.

Tracking Multiple Tables

The CDC Graph Generator operator can handle only a single table, so if you must track multiple tables, you must use the operator for each source-target pair. However, based on how much change data is generated in the source, you can optimize resource usage in the SAP Data Intelligence cluster instead of creating a graph for every table.

Example: 10 source tables, 2 are rapidly changing and 8 are slowly changing.

For rapidly changing tables, create one graph for each rapidly changing table. This graph runs indefinitely and polls the source for changes rapidly. In addition, it is recommended that you set *Delta Graph Run Mode* to *polling interval* with a short polling interval so that the operator fetches data rapidly as well.

For slowly changing tables, you can chain multiple CDC Graph Generator operators one after another in the same graph. While chaining the operators, set the *Delta Graph Run Mode* to Manual, so that the first operator completes its data movement and before the graph proceeds to the second one. While chaining, you can also group every N CDC Graph Generator operator in a single group so that they run in their own pod.

At the same time, to save further resource usage, schedule the graph to run periodically based on the delta change rate in the source. This way, you are not consuming resources for slowly changing tables at the same time, and are not using too many SQL connections to the source.

Recommendations

- Recommended group resource is 0.5 CPU and 1500 m memory. More than 2 CPU is not helpful for processing; however, more RAM (maximum 4 Gb) is better when there is a chain of CDC operators.
- RAM requirements are based on the `fetchSize` and Table Row Size.
 - For wide tables (for example, more than 100 columns), a smaller `fetchSize` (<1000) and higher RAM value is recommended.
 - For narrow tables, a larger `fetchSize` provides better throughput.

```
"groupResources": {
  "memory": {
    "request": "512M"
  },
  "cpu": {
    "request": "0.5"
  }
},
```

- During delta processing at maximum, you may use 2 SQL connections to the source table; however, for the initial load, use 1 + *Max Number of Loaders* connections.

Related Information

[Changing Data Capture \(CDC\)](#)
[CDC Graph Generator \(Deprecated\)](#)

10 Sizing for Metadata Explorer and Self-Service Data Preparation

Fine-tune performance of Metadata Explorer and Self-Service Data Preparation content using the App-Data tenant application. App-Data contains Metadata Explorer and Self-Service Data Preparation as nested applications.

App-Data is a tenant application that can scale to hundreds of users using a single pod, versus a user application that has one pod per user. Most systems run into resource issues before they reach 100 pods. We expose several configurations in App-Data to help fine-tune performance.

SAP tested Metadata Explorer and Self-Service Data Preparation for performance using default options. In the testing, SAP made the following assumption:

- A user makes a backend call every 6–10 seconds.
- Any tests with no time between calls have a conservative factor of 5X to scale up to the number of users supported.

The following table shows the test results.

Metadata Explorer	Self-Service Data Preparation
Tests show good performance from 1–100 users. A decrease in performance is noted between 100–200 users, with a maximum of 500 users tested on a single instance.	Tests show good performance from 1–50 users. A decrease in performance is noted between 50–100 users, with a maximum of 150 users tested on a single instance.


Related Information

[Configure App-Data \[page 264\]](#)

[Configure Other Applications \[page 266\]](#)

10.1 Configure App-Data

App-Data application settings are exposed to help you configure App-Data for optimal performance.

You can resolve HANA connection bottlenecks with HANA pool options in the App-Data application. Bottlenecks in other applications require changes to those applications and are listed in [Configure Other Applications \[page 266\]](#). For these changes to take effect, the daemon or app-data pods must be restarted. Restart the *Data Application* or *Data App Daemon* by clicking  *Restart*.

Data Application Options

Option	Description
HANA Pool Timeout	<p>The HANA pool timeout in milliseconds (500, 100000) for each nested application. The default value is 5000.</p> <p>The Metadata Explorer and Self-Service Data Preparation applications each have their own connection pool for accessing the HANA instance. The HANA Pool Timeout option is used to control the timeout value in both pools. The default of 5000 is good for most cases. If you encounter errors with the pool, increasing the value to 10,000–15,000 may keep the errors from occurring; however, it could result in longer wait times before the error is returned. The most common pool error logs the following message: <code>{"message": "could not acquire connection from pool", "module": "dh-app-metadata"}</code></p>
HANA Maximum Pool Size	<p>The HANA maximum pool size (1, 300) for each nested application. The default value is 100.</p> <p>The Metadata Explorer and Self-Service Data Preparation applications each have their own connection pool for accessing the HANA instance. The HANA Maximum Pool Size option is used to control the maximum number of connections in the pool, which corresponds to the number of concurrent database accesses that can be executed in parallel (each request to the backend gets a connection from the pool before it is processed). The default of 100 is good for most use cases. If requests fail to connect with the pool before the pool times out, you can increase the value. We recommend 100–150 connections. The most common pool error logs the following message: <code>{"message": "could not acquire connection from pool", "module": "dh-app-metadata"}</code></p>
Maximum number of concurrent tasks that execute pipelines in Modeler (automatic lineage excluded)	<p>Set to 10 or less to have up to 10 parallel batch processes running. Set to -1 to turn this option off.</p> <p>This process checks before each rulebook, preparation, profile, and publication that there are less than the maximum number of tasks are currently running.</p>
Time to live after logout	<p>Enter the number of minutes before the pod is stopped when no one is using metadata. It saves resources when the system is not being used for an extended amount of time.</p>

Data Application Daemon Options

Option	Description
HANA Pool Timeout	<p>The HANA pool timeout in milliseconds (500, 100000) for each nested application. The default value is 5000.</p> <p>The daemon pod that handles background processes also has a connection pool for accessing HANA. The daemon pod HANA Pool Timeout is used to control the timeout value of the pool. The default of 5000 is good for most cases. If you encounter errors with the pool, increasing the value to 10,000–15,000 may keep the errors from occurring.</p>
HANA Maximum Pool Size	<p>The HANA maximum pool size (1, 300) for each nested application. The default value is 50.</p> <p>The daemon pod that handles background processes also has a connection pool for accessing HANA. The daemon pod HANA Maximum Pool Size option is used to control the maximum number of connections in the pool. The default of 50 is good for most cases. If you encounter errors with the pool, increasing the value to 75–100 may fix the issue.</p>

10.2 Configure Other Applications

Other application settings are exposed to help you configure applications for optimal performance.

The Metadata Explorer and Preparation applications sit on top of the stack and call other applications that can potentially become bottlenecks. You can change the following application configurations to improve performance in other applications that are called by Metadata Explorer and Preparation.

Flowagent

Option	Description
Instances	<p>Browse remote connection, view dataset summary, view factsheet, and factsheet preview, rulebook dataset binding, and other calls may call the flowagent if that object is not published. If these types of calls become slower, the slowdown may be with flowagent. Increasing the number of flowagent instances may help.</p> <p>The default value is 1.</p>

11 Understanding Security

Learn about the security approach SAP Data Intelligence uses and about additional ways that you can increase security.

Related Information

[Data Protection and Privacy in SAP Data Intelligence \[page 267\]](#)

[Security Recommendations for SAP Data Intelligence \[page 272\]](#)

[Securing SAP Data Intelligence \[page 277\]](#)

11.1 Data Protection and Privacy in SAP Data Intelligence

SAP Data Intelligence provides the technical enablement and infrastructure to allow you to run applications on SAP Data Intelligence to conform to the legal requirements of data protection in the different scenarios in which SAP Data Intelligence is used.

Introduction to Data Protection

For general information about data protection and privacy in SAP BTP, see the SAP BTP documentation under [Data Protection and Privacy](#).

Data protection is associated with numerous legal requirements and privacy concerns. In addition to compliance with general data privacy acts, it is necessary to consider compliance with industry-specific legislation in different countries. This section describes the specific features and functions that SAP Data Intelligence provides to support compliance with the relevant legal requirements and data privacy.

This guide does not give advice on whether these features and functions are the best method to support company, industry, regional, or country-specific requirements. Furthermore, this guide does not give advice or recommendations about additional features that would be required in a particular environment. Decisions related to data protection must be made on a case-by-case basis and under consideration of the given system landscape and the applicable legal requirements.

Note

In most cases, compliance with data privacy laws is not a product feature. SAP software supports data privacy by providing security features and specific functions relevant to data protection, such as functions for the simplified blocking and deletion of personal data. SAP does not provide legal advice in any form. The definitions and other terms used in this guide are not taken from any given legal source.

Glossary

Term	Definition
Personal data	Information about an identified or identifiable natural person.
Business purpose	A legal, contractual, or in other form justified reason for the processing of personal data . The assumption is that any purpose has an end that is usually already defined when the purpose starts.
Blocking	A method of restricting access to data for which the primary business purpose has ended.
Deletion	Deletion of personal data so that the data is no longer usable.
Retention period	The time during which data must be available.
End of purpose (EoP)	A method of identifying the point in time for a data set when the processing of personal data is no longer required for the primary business purpose . After the EoP has been reached, the data is blocked and can only be accessed by users with special authorization.

SAP Data Intelligence Approach to Data Protection

Many data protection requirements depend on how the business semantics or context of the data stored and processed in SAP Data Intelligence are understood.

Note

Using capabilities to communicate with other data sources, SAP Data Intelligence may also be used to process data that is stored in other systems and accessed through virtual tables.

In SAP Data Intelligence installations, the business semantics of data are part of the application definition and implementation. SAP Data Intelligence does not own or store any sensitive data but the features for working with data sources, flowgraphs, and so on. Therefore, it is the user that knows, for example, which tables in the database contain sensitive personal data, or how business level objects, such as sales orders, are mapped to technical objects in the SAP Data Intelligence ecosystem.

Caution

SAP Data Intelligence trace and dump files may potentially expose personal data.

The following data protection and privacy functions enable a company to process personal data in a clear and compliant manner:

Function	Description
Erase personal data	SAP Data Intelligence needs the user ID, password, and (optionally) name of SAP Data Intelligence users for its operations. These are managed through the SAP Data Intelligence System Management interface and you can delete them, if necessary. For any external data source that SAP Data Intelligence accesses that may contain personal data, you should delete data using the appropriate method for that source, because SAP Data Intelligence does not know which objects store personal data.
Log changes to personal data	For the personal data that SAP Data Intelligence owns, all changes are logged in the audit logs, which can be accessed through SAP Data Intelligence diagnostics or the logs.
Information about data subjects	For the personal data that SAP Data Intelligence owns, users can see the contents of stored data in SAP Data Intelligence System Management by navigating to user management. All stored data except passwords is shown on this page.
Log read access to sensitive personal data	For the personal data that SAP Data Intelligence owns, any access is logged in the audit logs, which can be accessed through SAP Data Intelligence diagnostics or the logs.
Consent for personal data	SAP Data Intelligence does not store any data that is subject to consent for its own operations. If any external data source which SAP Data Intelligence can access stores any personal data, the owner of the data source must obtain the consent.

Note

Database trace and dump files may potentially expose personal data, for example, a trace set to a very high trace level, such as DEBUG or FINE.

SAP Data Intelligence provides a variety of security-related features to implement general security requirements that are also required for data protection and privacy:

Aspect of Data Protection and Privacy	More Information
Access control	Roles and scopes Enabling Authentication for SAP Data Intelligence Services and Users [page 278]
Transmission control/communication security	Data Provisioning Agent documentation <i>SAP Data Intelligence Self-Signed CA, X.509 Certificates and TLS for SAP Data Intelligence Services</i>

Aspect of Data Protection and Privacy	More Information
Separation by purpose	Roles and scopes Enabling Authentication for SAP Data Intelligence Services and Users [page 278]

⚠ Caution

The extent to which data protection is ensured depends on secure system operation. Network security, security note implementation, adequate logging of system changes, and appropriate usage of the system are the basic technical requirements for compliance with data privacy legislation and other legislation

Data Protection and Privacy in Graphs

For a threat model, the protection of data and ensurance of data privacy is a major concern. Graphs that are executed within the SAP Data Intelligence Modeler are configured, instructed, and started by customer request. Therefore, the system behaves as a data processor, whereas the user or customer of the system is the data controller. Hence, the system neither audit logs the input of personal or sensitive data from source systems, nor audit logs transformations or ingestions into target systems. The data owner or the owner of the source and target systems is responsible to ensure tracability and instruct the systems to properly generate relevant audit logs, which allows the customer to be compliant to local data protection and privacy laws.

Related Information

[Managing Audit Logs \[page 270\]](#)

[Viewing Audit Logs \[page 271\]](#)

[Malware Scanning \[page 272\]](#)

11.1.1 Managing Audit Logs

All SAP Data Intelligence components write audit logs for accessing, modifying, and erasing personal data, and editing your security configuration.

SAP Data Intelligence stores application audit logs that are customer-specific in the Audit Log Service. With the applicable permission, you can retrieve these audit logs using a tool such as the SAP Business Technology Platform (BTP) Audit Log viewer.

There are four types of audit logs:

- Configuration Change
- Personal Data Access
- Personal Data Modification

- Security Event

For more information about these audit logs, see [Audit Log Write API for Customers](#) in the SAP BTP documentation.

Note

The core SAP Data Intelligence system stores audit logs internally for debugging purposes only. You can't change persistence for these audit logs.

11.1.2 Viewing Audit Logs

SAP Data Intelligence provides a comprehensive audit logging system, which includes events that are related to Data Protection Principles.

The audit logs include both audit logs from SAP Data Intelligence applications and infrastructure audit logs of the SAP Data Intelligence cluster while using the SAP BTP Audit Log Services.

To see infrastructure audit logs of SAP Data Intelligence, subscribe to the Audit Log Viewer application instance in your subaccount. After you subscribe, the [audit log viewer](#) role is available from the trust configuration panel. Only a subaccount security administrator can assign the audit log viewer role to you from the trust configuration panel, which lets you view audit logs. The subaccount security administrator role is present only for the creator of the subaccount, or it can be transferred from one subaccount security administrator to another account.

Adding Another Subaccount Security Administrator

If you are a subaccount security administrator, perform the following steps to add another administrator.

1. In the [Security](#) tab, navigate to [Administrators](#).
2. Click the [Add Administrators](#) button, provide the [User ID](#), and click [OK](#).

The user should now have administrator privileges in the subaccount.

Granting Permissions to View Audit Logs

Perform the following steps to grant permission to view audit logs.

1. In the Subaccount page, navigate to [Role Collections](#) under the [Security](#) tab.
2. Click [New Role Collection](#), enter a name for the role collection, and click [Save](#). The new role collection appears in the role collection list.
3. Select the new role collection, add the [auditlog-viewer!](#) role, and select [Auditlog_Auditor](#) as the [Role Template and Role](#).
After you configure the role collection, you can assign users the role.
4. Navigate to the subaccount page and choose the [Security](#) tab. In [Trust Configuration](#), choose the identity provider that you use and make sure that the [Role Collection Assignment](#) page is open.
5. To assign the new role to your user, provide either the e-mail or user ID, depending on your identity provider settings.

Access Audit Log View

To access the Audit Log Viewer application, navigate to [Subscriptions](#) from the subaccount page, then click [Go To Application](#) under the [Audit Log Viewer](#) subscription.

For more information about enabling the audit log viewer in SAP BTP, see [Audit Log Viewer for the Cloud Foundry Environment](#).

11.1.3 Malware Scanning

SAP Data Intelligence scans all uploaded files for malware.

SAP Data Intelligence allows users to upload files to enable some component features, such as SAP Data Intelligence Modeler and Machine Learning operators. The files are stored in an SAP Data Intelligence cluster, and they may be read by other users, depending on their permissions scope.

As a second level of defense, and to comply with compliance controls, SAP Data Intelligence scans uploaded files for malware. In case of a positive finding, the upload fails and SAP Data Intelligence logs a security event.

The feature is enabled for all cloud instances.

11.2 Security Recommendations for SAP Data Intelligence

SAP Data Intelligence is delivered with secure default configurations wherever possible. However, you may want to review some settings and adjust them to your particular use case and corporate policies.

Security Recommendations

The following is a list of recommendations and their priority for you to review and adjust within your SAP Data Intelligence environment.

- For a description of each column in the table, see [Explanation of Table Headings \[page 277\]](#).
- For a description of the priorities, see [Explanation of Priorities \[page 277\]](#).

Priority	Secure Operations Map	Topic	Default Setting or Behavior	Recommendation	More Information	Index
Recommended	Client Security	Front-end Security: Browser	n/a	SAP recommends that all users maintain secure configurations and the latest patches for all Internet browsers.		SysMan.01
Critical	Roles and Authorization	Policies: Admin role	n/a	Restrict users with the administration role.	Manage Policies [page 50]	SysMan.02

Priority	Secure Operations Map	Topic	Default Setting or Behavior	Recommendation	More Information	Index
Recommended	User and Identity Management	Technical Users: Credentials	Passwords of technical users don't expire.	Rotate credentials of technical user regularly to restrict impact of leaked credentials.	Create and Delete users and User Instances [page 47]	SysMan.03
Recommended	User and Identity Management	Business User: SSO	SAP Data Intelligence integrates by default with the customers' BTP Extended Services for User Account and Authentication (XSUAA) instance of the related subaccount.	Configure a custom identity provider in BTP XSUAA to enable identity federation and allow single sign-on (SSO) for business users.	Configuring External Identity Providers in SAP Data Intelligence [page 278]	SysMan.04
Advanced	Roles and Authorization	Policies: Auto mapping	n/a	Define attribute mappings to automatically (de-)assign policies for business users (federated identities) based on role collections or other attributes managed in BTP XSUAA.	Manage Policies [page 50]	SysMan.05
Recommended	User and Identity Management	Technical Users: Authentication	Credential-based authentication.	Use client certificate authentication for internal technical users, assign complex passwords and rotate regularly, or disable password-based login for a user completely.	Create and Delete users and User Instances [page 47]	SysMan.07

Priority	Secure Operations Map	Topic	Default Setting or Behavior	Recommendation	More Information	Index
Recommended	Roles and Authorization	Policies	SAP delivers default policies for different personas to explore the functionality.	Don't use default compound policies or roles, but define and assign your own tailored policies and roles to enforce least privilege and separations of duty for your business users.	Manage Policies [page 50]	SysMan.08
Recommended	Roles and Authorization	Policies	SAP delivers a default policy that allows access to all connections.	Don't assign allowlist connection access policy in productive systems to your business users to ensure least privilege.	Manage Policies [page 50]	SysMan.09
Recommended	User and Identity Management	Technical Users: Reset on first login	n/a	Enforce a password reset for newly created technical users on the first login.		SysMan.10
Critical	Roles and Authorization	Business Users: Regular authorization review	n/a	Conduct a review of authorizations regularly to prevent privilege creep.	Manage Policies [page 50]	SysMan.11
Critical	Connected Systems	Security Monitoring and Forensics: Customer-controlled source systems	n/a	Enable audit logging in your connected systems that are configured via Connection Management and used by graphs, Jupyter Notebooks, or other objects to ensure authenticity and compliance to data protection and privacy regulations.	Data Protection and Privacy in SAP Data Intelligence [page 267]	CM.01

Priority	Secure Operations Map	Topic	Default Setting or Behavior	Recommendation	More Information	Index
Critical	Connected Systems	Customer-controlled source systems	n/a	Connected systems are the responsibility of customers (for example, regarding endpoints, channel and rest encryption, security configuration settings, authentication methods, general hardening, and so on).		CM.02
Recommended	Connected Systems	Restricted users to source systems	n/a	Use distinct technical users for connections to ensure separation of duties, least privilege, and authenticity.		CM.03
Recommended	Connected Systems	Connection User: Restrict privileges	n/a	Least privilege for technical connection users.		CM.04
Critical	Connected Systems	Connection Settings: Channel encryption	TLS if available for connection type.	Use TLS for channel encryption when possible.		CM.05
Recommended	Connected Systems	Connected Systems	Strongest method is pre-selected.	Use client certificate authentication or the most secure available authentication mode.		CM.06
Recommended	Connected Systems	Connection User: Technical credentials	n/a	Ensure regular rotation of connection credentials.		CM.07

Priority	Secure Operations Map	Topic	Default Setting or Behavior	Recommendation	More Information	Index
Advanced	Security Monitoring and Forensics	Security Audit Log: Review	Security and DPP-related events are audit logged.	Use BTP Audit Log Viewer in your subaccount to do a regular review or sampling of all customer audit logs.	Viewing Audit Logs [page 271]	SIEM-01
Advanced	Security Monitoring and Forensics	Security Audit Log: Deletion	Security and DPP-related events are automatically deleted.	Customer audit logs are stored for a restricted time period and then deleted.	Viewing Audit Logs [page 271]	SIEM-02
Recommended	Custom Code Security	Programming in SAP Data Intelligence Modeler	Consider the code security when you create graphs, operators, and Dockerfiles in the Modeler.	Don't add sensitive information into graph operators, Dockerfiles, and so on.	Creating Operators Creating Dockerfiles	DI-MOD-0001
Recommended	Security Hardening	Strong Authentication	Choose a connection type based on data residency requirements and applicable regulations. There is no default value of connection type. Choose actively.	Use secure connections and strong authentication types in graphs and operators.	Using Managed Connections in Script Operators	DI-MOD-0002
Recommended	Data Protection and Privacy	Encryption, Data in Transit	You can use a different transport protocol for data transfer in graphs.	Use only encrypted channels for graphs, such as TLS, and use the strongest available authentication type.		DI-MOD-0003

Explanation of Table Headings

Priority	Secure Operations Map	Topic	Default Setting or Behavior	Recommendation	More Information	Index
The priority is based on the risk related to the recommendation. For a description of the priorities, see Explanation of Priorities [page 277] .	The Secure Operations Map is a reference model to structure the broad area of security for content and discussions. For more information about the Secure Operations Map, see SAP Security Optimization Services Portfolio .	The topic is a categorization to find similar topics across services.	Description of the default state or behavior.	The SAP recommendation for the configuration.	A link to documentation that explains how you can achieve the recommendation.	A stable reference to identify the recommendation.

Explanation of Priorities

Priority	Description
Critical	Exposes significant risk or threatens system reliability.
Recommended	Improves the security of the landscape and significantly reduces the attack surface.
Advanced	Extends the recommendation to a higher standard, addresses areas where there's no clear good or bad, and addresses your organization-specific requirements.

11.3 Securing SAP Data Intelligence

When using a distributed system, you must be sure that your data and processes support your business needs without allowing unauthorized access to critical information. User errors, negligence, or attempted manipulation of your system should not result in loss of information or processing time.

These demands on security apply likewise to SAP Data Intelligence.

Related Information

[Enabling Authentication for SAP Data Intelligence Services and Users \[page 278\]](#)

[Configuring External Identity Providers in SAP Data Intelligence \[page 278\]](#)

[Giving User Permissions for SAP Data Intelligence Access \[page 279\]](#)

[Connecting Your On-Premise Systems To SAP Data Intelligence \[page 280\]](#)

11.3.1 Enabling Authentication for SAP Data Intelligence Services and Users

SAP Data Intelligence requires authentication for all users and services.

SAP Data Intelligence supports the following authentication mechanisms for all services:

- User name and password authentication.
- User client certificate authentication.

After an administrator creates a user, the user becomes enabled on all SAP Data Intelligence endpoints, including all user interfaces and programmatic endpoints.

Note

Synchronizing a user for all SAP Data Intelligence services can take up to 5 minutes.

The SAP Data Intelligence installer creates the first user who is the default tenant. If you have the correct permissions, you can create and manage additional users through the SAP Data Intelligence System Management application user management module. The SAP Data Intelligence installer also creates service users for internal authentication between SAP Data Intelligence services.

The following table summarizes the SAP Data Intelligence endpoints that are available to users and other services (exposed from Kubernetes by default) and the authentication methods that apply to them.

Endpoint	Authentication Type and Proposed Security Measures
System Management application user interface	SAP Data Intelligence users with a user interface authentication form.
System Management REST endpoints	User client certificate authentication.

Related Information

[Create and Delete users and User Instances \[page 47\]](#)

[Manage Certificates \[page 122\]](#)

11.3.2 Configuring External Identity Providers in SAP Data Intelligence

SAP Data Intelligence service integrates with SAP BTP User Account and Authentication (SAP BTP UAA). This allows the federation of external and custom identity providers (IdP). In that context, SAP BTP UAA works as a

service broker. For example, it allows the configuration of identity providers such as Security Assertion Markup Language (SAML).

Configuring External Identity Providers

More information is available in the SAP BTP documentation:

- For information about the configuration of identity providers via SAP BTP UAA, refer to [Data Privacy and Security](#) in the SAP BTP documentation.
- For information about establishing trust between SAP BTP UAA and SAP BTP Identity Authentication Service (IAS), refer to [Establish Trust and Federation with UAA Using SAP BTP Identity Authentication Service](#).

The configuration is handled within the SAP BTP security configuration dashboard of the subaccount that contains the SAP Data Intelligence instance.

Every user who can be federated via the SAP BTP UAA can log into SAP Data Intelligence with the initial role of a member.

After the first login, broader access can be granted to the externally federated user.

For more information about user handling and access control settings, see [SAP BTP Administration](#).

SAP Data Intelligence stores the e-mail addresses of users from external identity providers. SAP Data Intelligence uses e-mail addresses for preregistering users. E-mail addresses serve as display names. The user information, including e-mail addresses, is readable by logged-in users from the same tenant through SAP Data Intelligence APIs.

Related Information

SAP Vora Integration for External Users

11.3.3 Giving User Permissions for SAP Data Intelligence Access

To authenticate an external user for an SAP Data Intelligence service, you must assign the `sap.dh.systemAccess` policy to the user in SAP Data Intelligence System Management.

After the external user logs in for the first time, the tenant administrator must navigate to the [Users](#) page in [System Management](#) and assign one of the policies that nests, or references, `sap.dh.systemAccess`. For more information about policy assignment, see [Manage Policies \[page 50\]](#).

External users with the `sap.dh.systemAccess` policy can log into the SAP Data Intelligence service by clicking the [SAP CP XSUAA](#) button on the login page.

11.3.4 Connecting Your On-Premise Systems To SAP Data Intelligence

To connect your on-premise systems to SAP Data Intelligence, submit a support request. The SAP operations team will help you set it up.

There are three connectivity options available to connect to SAP Data Intelligence:

- SAP Cloud Connector
- Site-to-Site VPN
- VPC Peering

The approach that you use depends on whether your on-premise environment is hosted. We recommend that you use SAP Cloud Connector, which serves as a link between SAP Business Technology Platform (BTP) applications, such as SAP Data Intelligence and on-premise systems. If you use one of the following hosts and SAP Cloud Connector is not suitable for your scenario, use a peering connection:

Hosted In	Use
Amazon Web Services (AWS)	AWS Virtual Private Cloud (VPC) peering connection
Microsoft Azure	VNet peering

If your on-premise environment is not hosted, a Virtual Private Network (VPN) using the AWS Site-to-Site VPN is required.

Related Information

[Connect Using Site-to-Site VPN \[page 280\]](#)

[Connect Using Virtual Network Peering \[page 282\]](#)


[Configure Cloud Connector \[page 22\]](#)

11.3.4.1 Connect Using Site-to-Site VPN

If your on-premise environment is not hosted in Amazon Web Services (AWS) or Microsoft Azure, which are clouds where the SAP Data Intelligence can be hosted, connect it to SAP Data Intelligence with a Virtual Private Network (VPN).

Context

VPN connectivity provides an added layer of security and easier traffic routing between your on-premise systems and SAP Data Intelligence. You can connect to SAP Data Intelligence through the Internet or through an IP security (IPSec) VPN tunnel.



Internet connectivity to SAP Data Intelligence APIs and user interfaces using secure protocols are always available. To add another network security layer, you can establish a VPN connection between your on-premise or other hosted environment and SAP Data Intelligence. Submit a support request to the SAP operations team, and the connection will be jointly implemented as follows:

Procedure

1. Contact your network department to verify that they will support a VPN connection to SAP Data Intelligence.
2. Follow the steps described in SAP Note [3108686](#).

Restrictions

Note the following restrictions:

- To ensure that your IPsec VPN endpoint is compatible, see <https://docs.aws.amazon.com/vpc/latest/adminguide/Welcome.html> or <https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-about-vpn-devices>, depending where your SAP Data Intelligence instance is hosted.
- VPN provides a transparent connection to your network. The SAP Data Intelligence internal default Classless Inter-Domain Routing (CIDR) block is 10.0.0.0/16. The IP address range cannot be changed after the SAP Data Intelligence tenant landscape has been created. Therefore, ensure that you do not have an IP address range conflict with the SAP Data Intelligence default network. If the default CIDR segment conflicts with your network, specify an appropriate non-overlapping CIDR block (/22 or larger) when you create the SAP Data Intelligence service instance. For more information, see [Create an SAP Data Intelligence Cloud Instance in SAP BTP \[page 8\]](#).
- High availability setup of a VPN connection is not available.
- Because a VPN connection leverages the Internet, SAP cannot provide performance and availability service-level agreements. Performance depends on your on-premise Internet Service Provider connectivity capacity and other factors.
- If your corporate hosting is also on Amazon AWS, a VPN connection to SAP Data Intelligence hosted on AWS is not supported by Amazon. Contact SAP support to discuss alternatives.

11.3.4.2 Connect Using Virtual Network Peering

When your on-premise environment is hosted by the virtual network for Amazon Web Service or Azure, use the virtual network's peering connection to connect your on-premise environment to SAP Data Intelligence cloud.

Context

SAP Data Intelligence supports the virtual networks listed in the following table with the supported peering method.

Virtual Network	Virtual Network Peering
Amazon Web Service (AWS)	VPC (Virtual Private Cloud) peering
Azure	VNet

Both VPC and VNet peering create network connectivity between two virtual networks that are owned by different account holders. For example, VPC peering creates network connectivity between two AWS VPCs that are owned by different AWS account holders. To establish virtual network peering between your hosted environment and SAP Data Intelligence, perform the following steps:

Procedure

1. Contact your network department to confirm that your virtual network peering method is compliant with your organization rules.
2. Follow the steps described in SAP Note [3108686](#) that walk you through the following processes:
 - VPN site-to-site configurations for AWS and Azure.
 - Virtual network peering.

Restrictions

Note the following restrictions:

- VPC peering provides a transparent connection to your network. The SAP Data Intelligence internal default Classless Inter-Domain Routing (CIDR) block is 10.0.0.0/16.
- The IP address range can't be changed after the SAP Data Intelligence tenant landscape is created. Therefore, ensure that you don't have an IP address range conflict with the SAP Data Intelligence default network.
- If the default CIDR segment conflicts with your network, specify an appropriate non-overlapping CIDR block (/22 or larger) when you create the SAP Data Intelligence service instance. For more information, see [Create an SAP Data Intelligence Cloud Instance in SAP BTP \[page 8\]](#).

Find more information about common issues in VPN configuration in the SAP Note [3108686](#).

12 Troubleshooting SAP Data Intelligence

Contains information that helps you troubleshoot problems in SAP Data Intelligence.

Related Information

[Troubleshooting SAP Cloud Connector \[page 284\]](#)

[Troubleshooting Flowagent \[page 285\]](#)

12.1 Troubleshooting SAP Cloud Connector

Contains information that helps you troubleshoot problems when using SAP Cloud Connector with SAP Data Intelligence.

Failed Connection

Symptom

Connection fails to be established when trying to use SAP Cloud Connector.

Analysis

Checking logs of SAP Cloud Connector (on customer side) reveals an error message like the following: Caused by: sun.security.validator.ValidatorException: PKIX path building failed: sun.security.provider.certpath.SunCertPathBuilderException: unable to find valid certification path to requested target

Root cause

The certificate of the SAP Data Intelligence ingress is not trusted by SAP Cloud Connector.

Solution

Obtain the certificate from the SAP Data Intelligence ingress endpoint and upload it into the trust store of the SAP Cloud Connector. For an example, see [Import the Git Server Certificate into the JVM](#).

12.2 Troubleshooting Flowagent

Contains information that helps you troubleshoot Flowagent service issues in SAP Data Intelligence.

Flowagent service can hang indefinitely for database connections

There are some situations where registering an invalid connection leads to a hang in the Flowagent service, which is responsible for checking status, browsing and viewing metadata of connections. This issue can occur, for example, if you register a Microsoft SQL Server connection using MySQL information; however, there are other invalid combinations that can lead to the same behavior.

The connections affected by this issue may include: Azure SQL DB, DB2, MSSQL, MySQL, Oracle, and Redshift.

Workaround



In SAP Data Intelligence, restart the Flowagent application via SAP Data Intelligence System Management.

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.

© 2026 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.

