



Informatica® Data Ingestion and Replication
July 2024

Connectors and Connections

© Copyright Informatica LLC 2019, 2024

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC.

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation is subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License.

Informatica, Informatica Cloud, Informatica Intelligent Cloud Services, PowerCenter, PowerExchange, and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties. Required third party notices are included with the product.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, report them to us at infa_documentation@informatica.com.

Informatica products are warranted according to the terms and conditions of the agreements under which they are provided. INFORMATICA PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2024-07-12

Table of Contents

Preface	6
Informatica Resources.	6
Informatica Documentation.	6
Informatica Intelligent Cloud Services web site.	6
Informatica Intelligent Cloud Services Communities.	6
Informatica Intelligent Cloud Services Marketplace.	7
Data Integration connector documentation.	7
Informatica Knowledge Base.	7
Informatica Intelligent Cloud Services Trust Center.	7
Informatica Global Customer Support.	7
 Chapter 1: Connectors and Connections.....	8
 Chapter 2: Data Ingestion and Replication connectors.....	9
Application Ingestion and Replication connectors.	9
Database Ingestion and Replication connectors.	11
Mock connectors.	13
File Ingestion and Replication connectors.	14
Streaming Ingestion and Replication connectors.	15
 Chapter 3: Data Ingestion and Replication connection properties.....	17
Configuring a connection.	17
Adobe Analytics Mass Ingestion connection properties.	18
Advanced FTP V2 connection properties.	19
Advanced FTPS V2 connection properties.	21
Advanced SFTP V2 connection properties.	23
Amazon Kinesis connection properties.	25
Amazon Kinesis Firehose connection properties.	25
Amazon Kinesis Streams connection properties.	27
AWS Credential Profile.	28
Amazon Redshift V2 connection properties.	28
Prepare for authentication.	28
Create a minimal Amazon IAM policy.	30
Configure IAM authentication.	31
Configure an assume role for Amazon Redshift.	31
Configure an assume role for Amazon S3 staging.	34
Enable encryption.	37
Connect to Amazon Redshift.	38
Proxy server settings.	48
Private communication with Amazon Redshift.	48

Amazon S3 V2 connection properties.	49
Credential Profile File Authentication.	53
Private communication with Amazon S3.	54
AMQP connection properties.	54
Business 360 Events connection properties.	56
Cloud Integration Hub connection properties.	56
Databricks connection properties.	57
Staging prerequisites.	57
SQL warehouse.	57
Databricks cluster.	59
Connect to Databricks.	60
JDBC URL parameters.	66
Rules and guidelines for personal staging location.	66
Db2 for i Database Ingestion connection properties.	67
Db2 for LUW Database Ingestion connection properties.	68
Db2 for zOS Database Ingestion connection properties.	69
Flat file connection properties.	70
Google Analytics Mass Ingestion connection properties.	73
Google BigQuery V2 connection properties.	73
Google Cloud Storage V2 connection properties.	75
Configuring the proxy settings on Windows.	76
Configuring the proxy settings on Linux.	77
Google PubSub - Streaming Ingestion and Replication connection properties.	79
Hadoop Files V2 connection properties.	80
JDBC V2 connection properties.	81
JMS connection properties.	83
Kafka connection properties.	84
Configuring the krb5.conf file to read data from or write to a Kerberised Kafka cluster.	87
Configuring SASL PLAIN authentication for a Kafka cluster.	89
Configuring SASL_SSL authentication for a Cloud Confluent Kafka cluster.	90
Connecting to Amazon Managed Streaming for Apache Kafka.	91
Marketo V3 connection properties.	92
Microsoft Azure Blob Storage V3 connection properties.	93
Microsoft Azure Data Lake Storage Gen2 connection properties	94
Microsoft Azure Event Hub connection properties.	96
Microsoft Azure Synapse Analytics Database Ingestion connection properties.	97
Microsoft Azure Synapse SQL connection properties.	98
Microsoft Dynamics 365 Mass Ingestion connection properties.	100
Microsoft SQL Server connection properties.	104
Microsoft Fabric OneLake connection properties.	106
MongoDB Mass Ingestion connection properties.	107
MQTT connection properties.	108

MySQL connection properties.	110
Netezza connection properties.	111
NetSuite Mass Ingestion connection properties.	112
OPC UA connection properties.	114
Oracle Cloud Object Storage connection properties.	116
Oracle Database Ingestion connection properties.	117
Oracle Fusion Cloud Mass Ingestion connection properties.	123
PostgreSQL connection properties.	124
REST V2 connection properties.	126
Salesforce Marketing Cloud connection properties.	136
Salesforce Mass Ingestion connection properties.	138
SAP HANA Database Ingestion connection properties.	140
SAP Mass Ingestion connection properties.	143
SAP ODP Extractor connection properties.	149
ServiceNow Mass Ingestion connection properties.	154
Snowflake Data Cloud connection properties.	155
Standard authentication.	156
OAuth 2.0 authorization code authentication.	157
Key pair authentication.	159
Set JDBC URL Parameters.	160
Private links to access Snowflake.	160
Teradata connection properties.	160
Workday Mass Ingestion connection properties.	162
Zendesk Mass Ingestion connection properties.	164
Index.	166

Preface

Use *Data Ingestion and Replication Connectors and Connections* to determine the types of connectors that you need to download to be able to access sources and targets for each type of ingestion and replication task. It also describes the properties that you configure when you define a connection to be used by a task. You can download connectors and define connections in Administrator.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit <https://docs.informatica.com>.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Intelligent Cloud Services web site

You can access the Informatica Intelligent Cloud Services web site at <http://www.informatica.com/cloud>. This site contains information about Informatica Cloud integration services.

Informatica Intelligent Cloud Services Communities

Use the Informatica Intelligent Cloud Services Community to discuss and resolve technical issues. You can also find technical tips, documentation updates, and answers to frequently asked questions.

Access the Informatica Intelligent Cloud Services Community at:

<https://network.informatica.com/community/informatica-network/products/cloud-integration>

Developers can learn more and share tips at the Cloud Developer community:

<https://network.informatica.com/community/informatica-network/products/cloud-integration/cloud-developers>

Informatica Intelligent Cloud Services Marketplace

Visit the Informatica Marketplace to try and buy Data Integration Connectors, templates, and mapplets:

<https://marketplace.informatica.com/>

Data Integration connector documentation

You can access documentation for Data Integration Connectors at the Documentation Portal. To explore the Documentation Portal, visit <https://docs.informatica.com>.

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit <https://search.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Intelligent Cloud Services Trust Center

The Informatica Intelligent Cloud Services Trust Center provides information about Informatica security policies and real-time system availability.

You can access the trust center at <https://www.informatica.com/trust-center.html>.

Subscribe to the Informatica Intelligent Cloud Services Trust Center to receive upgrade, maintenance, and incident notifications. The [Informatica Intelligent Cloud Services Status](#) page displays the production status of all the Informatica cloud products. All maintenance updates are posted to this page, and during an outage, it will have the most current information. To ensure you are notified of updates and outages, you can subscribe to receive updates for a single component or all Informatica Intelligent Cloud Services components. Subscribing to all components is the best way to be certain you never miss an update.

To subscribe, on the [Informatica Intelligent Cloud Services Status](#) page, click **SUBSCRIBE TO UPDATES**. You can choose to receive notifications sent as emails, SMS text messages, webhooks, RSS feeds, or any combination of the four.

Informatica Global Customer Support

You can contact a Global Support Center through the Informatica Network or by telephone.

To find online support resources on the Informatica Network, click **Contact Support** in the Informatica Intelligent Cloud Services Help menu to go to the **Cloud Support** page. The **Cloud Support** page includes system status information and community discussions. Log in to Informatica Network and click **Need Help** to find additional resources and to contact Informatica Global Customer Support through email.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at <https://www.informatica.com/services-and-training/support-services/contact-us.html>.

CHAPTER 1

Connectors and Connections

Connections provide access to data in cloud and on-premises applications, platforms, databases, and flat files. Before you can define a connection, ensure that the connector for the source or target type is installed in Informatica Intelligent Cloud Services.

If multiple connectors are available for a source or target type, get the one that your ingestion type supports. Some connectors are pre-installed. If you need a connector that is not pre-installed, you can download it from the **Add-On Connectors** page in Administrator.

CHAPTER 2

Data Ingestion and Replication connectors

You must have the correct connectors to create connections for the sources and targets you use in your Data Ingestion and Replication tasks.

Before you can define connections, your organization administrator must ensure that the source and target connectors that the organization uses are installed. Also, you must enable connectors for your runtime environment.

For more information about connectors and connections, see "Licenses," "Runtime Environment," and "Connections" in the Administrator help.

Application Ingestion and Replication connectors

Before you define a connection for application ingestion and replication tasks, verify that the connectors for your source and target types are available in Informatica Intelligent Cloud Services.

The following table lists the connectors that Application Ingestion and Replication requires to connect to a source or target:

Source or target type	Connector	Use for
Amazon Aurora PostgreSQL	PostgreSQL	Targets in initial load, incremental load, and combined initial and incremental load operations. Applies only for Salesforce source.
Adobe Analytics	Adobe Analytics Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations
Amazon Redshift	Amazon Redshift V2	Targets in initial load, incremental load, and combined initial and incremental load operations
Amazon S3	Amazon S3 V2	Targets in initial load, incremental load, and combined initial and incremental load operations
Apache Kafka	Kafka	Targets in incremental load operations
Databricks	Databricks	Targets in initial load, incremental load, and combined initial and incremental load operations

Source or target type	Connector	Use for
Google Analytics	Google Analytics Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations
Google BigQuery	Google BigQuery V2	Targets in initial load, incremental load, and combined initial and incremental load operations
Google Cloud Storage	Google Cloud Storage V2	Targets in initial load, incremental load, and combined initial and incremental load operations
Marketo	Marketo V3	Sources in initial load, incremental load, and combined initial and incremental load operations Note: Apache Kafka target supports only incremental load.
Microsoft Azure Data Lake Storage Gen2	Microsoft Azure Data Lake Storage Gen2	Targets in initial load, incremental load, and combined initial and incremental load operations
Microsoft Azure Synapse Analytics ¹	Microsoft Azure Synapse Analytics Database Ingestion	Targets in initial load, incremental load, and combined initial and incremental load operations
Microsoft Azure SQL Database	SQL Server	Targets in initial load operations. Applies only for Salesforce source.
Microsoft Fabric OneLake	Microsoft Fabric OneLake	Targets in initial load, incremental load, and initial and incremental load operations
Microsoft Dynamics 365	Microsoft Dynamics 365 Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations
NetSuite	NetSuite Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations
Oracle Fusion Cloud Applications	Oracle Fusion Cloud Mass Ingestion	<ul style="list-style-type: none"> - REST: Sources in initial load, incremental load, and combined initial and incremental load operations - BICC: Sources in initial load, incremental load, and combined initial and incremental load operations
Salesforce	Salesforce Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations
Salesforce Marketing Cloud	Salesforce Marketing Cloud	Sources in initial load operations
SAP ECC	<ul style="list-style-type: none"> - SAP Mass Ingestion - SAP ODP Extractor 	<ul style="list-style-type: none"> - SAP Mass Ingestion: Sources in initial load and incremental load, and combined initial and incremental load operations . - SAP ODP Extractor: Sources in initial load, incremental load, and combined initial and incremental load operations
SAP S/4HANA	SAP ODP Extractor	Sources in initial load, incremental load, and combined initial and incremental load operations
ServiceNow	ServiceNow Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations

Source or target type	Connector	Use for
Snowflake	Snowflake Data Cloud	Targets in initial load, incremental load, and combined initial and incremental load operations
Workday	Workday Mass Ingestion	<ul style="list-style-type: none"> - SOAP: Sources in initial load, incremental load, and combined initial and incremental load operations Note: Apache Kafka target supports only incremental load. - RaaS: Sources in initial load
Zendesk	Zendesk Mass Ingestion	Sources in initial load, incremental load, and combined initial and incremental load operations
1. For Microsoft Azure Synapse Analytics targets, Application Ingestion and Replication uses Microsoft Azure SQL Data Lake Storage Gen2 to store staging files. Before you configure a connection for a Microsoft Azure Synapse Analytics target, ensure that you have Microsoft Azure SQL Data Lake Storage Gen2 installed.		

Database Ingestion and Replication connectors

Before you begin defining connections for database ingestion and replication tasks, verify that the connectors for your source and target types are available in Informatica Intelligent Cloud Services.

The following table lists the connectors that Database Ingestion and Replication requires to connect to a source or target that can be configured in a database ingestion and replication task:

Source or target type	Connector	Use for
Amazon Redshift	Amazon Redshift V2	Targets in initial load, incremental load, and initial and incremental load jobs
Amazon S3	Amazon S3 V2	Targets in initial load and incremental load jobs
Databricks	Databricks	Targets in initial load, incremental load, and initial and incremental load jobs
Db2 for i	Db2 for i Database Ingestion	Sources in initial load, incremental load, and initial and incremental load jobs
Db2 for Linux, UNIX, and Windows	Db2 for LUW Database Ingestion	Sources in initial load jobs
Db2 for z/OS	Db2 for zOS Database Ingestion	Sources in initial load and incremental load jobs
Flat file	No connector required	Targets in initial load jobs
Google BigQuery	Google BigQuery V2	Targets in initial load, incremental load, and initial and incremental load jobs
Google Cloud Storage	Google Cloud Storage V2	Targets in initial load and incremental load jobs

Source or target type	Connector	Use for
Kafka, including Apache Kafka, Confluent Kafka, Amazon Managed Streaming for Apache Kafka, and Kafka-enabled Azure Event Hubs	Kafka	Targets in incremental load jobs
Microsoft Azure Data Lake Storage Gen2	Microsoft Azure Data Lake Storage Gen2	Targets in initial load and incremental load jobs
Microsoft SQL Server, including on-premises SQL Server, RDS for SQL Server, Azure SQL Database, and Azure SQL Managed Instance	SQL Server	Sources in initial load, incremental load, and combined initial and incremental load jobs. For Azure SQL Database sources, you must use the Query-based or CDC Tables capture method for incremental load and combined load jobs. Targets in initial load, incremental load, and initial and incremental load jobs.
Microsoft Azure Synapse Analytics ¹	Microsoft Azure Synapse Analytics Database Ingestion	Targets in initial load, incremental load, and initial and incremental load jobs
Microsoft Fabric OneLake	Microsoft Fabric OneLake	Targets in initial load, incremental load, and initial and incremental load jobs
MongoDB	MongoDB Mass Ingestion	Sources in initial load and incremental load jobs
MySQL, including RDS for MySQL	MySQL	Sources in initial load and incremental load jobs. RDS for MySQL in initial load jobs only.
Netezza	Netezza	Sources in initial load jobs
Oracle, including RDS for Oracle	Oracle Database Ingestion	Sources in initial load, incremental load, and initial and incremental load jobs Targets in initial load, incremental load, and initial and incremental load jobs
Oracle Cloud Infrastructure (OCI) Object Storage	Oracle Cloud Object Storage	Targets in initial load, incremental load, and initial and incremental load jobs
PostgreSQL, including on-premises PostgreSQL, Amazon Aurora PostgreSQL, Azure Database for PostgreSQL - Flexible Server, RDS for PostgreSQL, and Cloud SQL for PostgreSQL	PostgreSQL	Sources in initial load, incremental load, and initial and incremental load jobs Targets in initial load, incremental load, and initial and incremental load jobs (Amazon Aurora PostgreSQL only)
SAP HANA, including on-premises SAP HANA and SAP HANA Cloud	SAP HANA Database Ingestion	Sources in initial load and incremental load jobs
Snowflake	Snowflake Data Cloud	Targets in initial load, incremental load, and initial and incremental load jobs

Source or target type	Connector	Use for
Teradata Data Warehouse Appliance	Teradata	Sources in initial load jobs
1. For the Microsoft Azure Synapse Analytics target type, Database Ingestion and Replication uses Microsoft Azure SQL Data Lake Storage Gen2 to store staging files. Ensure that you have Microsoft Azure SQL Data Lake Storage Gen2 installed.		

Mock connectors

Database Ingestion and Replication supports mock, or sample, connections for some of the sources and targets. Use mock connections to learn how to create database ingestion and replication initial load tasks without creating real connections to the database.

A mock connector does not connect to a real database. Instead, a source mock connector uses flat files with sample data. A target mock connector reports the information about processed source data to Database Ingestion and Replication user interface, but it does not write any data to the target.

The sample connections appear in the source and target connection lists in Database Ingestion and Replication if you have the MockConnector license.

The following table lists mock connections that you can use for Database Ingestion and Replication sources and targets:

Connection name	Source or Target
Sample Oracle Connection	Source
Sample SQL Server Connection	Source
Sample S3 Connection	Target
Sample ADLS Gen2 Connection	Target

Note: You must use sample connections for both the source and target databases. You cannot use a sample connection for only one of them, for example, for the source but not for the target.

Source data

The source data for sample connections is stored in CVS files in the following directory:

```
Secure_Agent_installation/downloads/package-MockConnector.version/package/sampleData/  
source/database_type/
```

Each file represents a single table. A mock table name matches the file name. The first line in a file determines column headers, and the subsequent lines contain row data.

File Ingestion and Replication connectors

Before you define connections for file ingestion and replication tasks, ensure that you have a license for the connectors that File Ingestion and Replication requires for the source and target types.

The following table lists the connectors that a file ingestion and replication task supports based on the source and target types:

Source or target name	Connector	Source or target type
Local folder	No connector required	Source and Target
Advanced FTP	Advanced FTP V2 (add-on)	Source and Target
Advanced FTPS	Advanced FTPS V2 (add-on)	Source and Target
Advanced SFTP	Advanced SFTP V2 (add-on)	Source and Target
Amazon S3	Amazon S3 V2 (add-on)	Source and Target
Amazon Redshift	Amazon Redshift V2 (add-on)	target
Cloud Integration Hub	Cloud Integration Hub (add-on)	Source and Target Note: You can't configure a file ingestion and replication task using Cloud Integration Hub as both source and target.
Databricks	Databricks	Source and Target
Google BigQuery	Google BigQuery V2 (add-on)	Target
Google Cloud Storage	Google Cloud Storage V2 (add-on)	Source and Target
Hadoop Files	Hadoop Files V2 (add-on)	Source and Target
Microsoft Azure Blob Storage	Microsoft Azure Blob Storage V3 (add-on)	Source and Target
Microsoft Azure Data Lake Store	Microsoft Azure Data Lake Store Gen2 (add-on)	Source and Target
Microsoft Azure Data Lake Store	Microsoft Azure Data Lake Store V3 (add-on)	Source and Target
Microsoft Azure Synapse SQL	Microsoft Azure Synapse SQL (add-on)	Target
Microsoft Fabric OneLake	Microsoft Fabric OneLake (add-on)	Source and Target
Snowflake	Snowflake Cloud Data Warehouse V2 (add-on)	Target

Streaming Ingestion and Replication connectors

Before you define connections for the streaming ingestion and replication tasks, ensure that you have a license for the required connectors for your source and target types.

The following table lists the connectors that a streaming ingestion and replication task supports based on the source and target type:

Source or target name	Connector	Source or target type
Amazon Kinesis Data Firehose	Amazon Kinesis (add-on)	Target
Amazon Kinesis Data Streams	Amazon Kinesis (add-on)	Source and Target
Amazon Managed Streaming for Apache Kafka (Amazon MSK)	Kafka (add-on)	Source and Target
Amazon S3	Amazon S3 V2 (add-on)	Target
AMQP	AMQP (add-on)	Source
Apache Kafka	Kafka (add-on)	Source and Target
Azure Event Hubs Kafka	Kafka (add-on)	Source
Business 360 Events	Business 360 Events (add-on)	Source
Confluent Kafka	Kafka (add-on)	Source and Target
Databricks	Databricks	Target
Flat file	No connector required	Source and Target
Google BigQuery V2	Google BigQuery V2 (add-on)	Target
Google Cloud Storage	Google Cloud Storage V2 (add-on)	Target
Google PubSub	Google PubSub (add-on)	Source and Target
JDBC V2	JDBC V2 (add-on)	Target
JMS	JMS (add-on)	Source
Microsoft Azure Data Lake Storage	Azure Data Lake Store Gen2 (add-on)	Target
Microsoft Azure Event Hub	Azure Event Hubs (add-on)	Target
MQTT	MQTT (add-on)	Source
OPC UA	OPCUA (add-on)	Source
REST V2	REST V2 (add-on)	Source

Note: While importing a streaming ingestion and replication task, both read and write connection types appear in the drop-down list on the **Import Review** page. You can also see connections to connectors that are not supported by Streaming Ingestion and Replication.

CHAPTER 3

Data Ingestion and Replication connection properties

Connections provide access to data in cloud and on-premises applications, platforms, databases, and flat files. Connection definitions include the location of the source or target, the runtime environment, and the other properties specific to the connection type.

Before you can create a connection, ensure that the correct connectors for your sources and targets are available in Informatica Intelligent Cloud Services. The supported connectors vary by type of ingestion task.

To create a connection or search for an existing connection, use the Administrator service.

After you configure connection properties, the connection becomes available for use within the organization.

Configuring a connection

Configure a source or target connection on the **Connections** page in Administrator.

1. In Administrator, select **Connections**.
2. On the **Connections** page, click **New Connection**.
3. Configure the following connection details:

Property	Description
Connection Name	The name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Maximum length is 100 characters. Connection names are not case sensitive.
Description	A description of the connection.
Type	The type of connection, such as Amazon S3.

After you select the connection type, additional properties that are specific to that type appear.

4. Configure the connection-specific properties.
For example, if you are configuring an Amazon S3 connection, enter the Amazon S3 connection properties. Click the Help icon for a description of each connection property.

5. To test the connection, click **Test Connection**.

The results of the test are displayed at the top of the page.

If a connection fails, contact the database administrator, or recheck your settings and verify that the selected runtime environment has the status of Up and Running.

6. Click **Save** to save the connection.

Adobe Analytics Mass Ingestion connection properties

When you set up an Adobe Analytics Mass Ingestion connection, you must configure the connection properties.

Adobe Analytics uses a JSON Web Token (JWT) to authenticate the Adobe Analytics Mass Ingestion connection. To use an Adobe Analytics Mass Ingestion connection, you must create a Service Account Integration on Adobe Developer Console and then specify the service integration details in the connection properties. For more information about creating a Service Account Integration on Adobe Developer Console, see the [Adobe documentation](#).

The following table describes the connection properties for an Adobe Analytics Mass Ingestion connection:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Select the Adobe Analytics Mass Ingestion connection type.
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Client ID	Client ID of the Service Account that you created on Adobe Developer Console.
Client Secret	Client secret of the Service Account that you created on Adobe Developer Console.
Technical Account ID	Technical account ID of the Service Account.
Organization ID	Organization ID of the Service Account.
Private Key	Private key that is generated when you create the Service Account Integration. The private key is required to generate the JWT.

Connection property	Description
IMS Host	Base URL of Adobe Identity Management System (IMS). The default value is: <code>ims-na1.adobelogin.com</code>
IMS Exchange	Exchange URL of IMS. The connection use the JWT to obtain an access token from Adobe by making a POST request to the exchange URL. The default value is: <code>https://ims-na1.adobelogin.com/ims/exchange/jwt</code>

Advanced FTP V2 connection properties

When you set up an Advanced FTP V2 connection, you must configure the connection properties.

The following table describes the Advanced FTP V2 connection properties:

Connection property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Type	Select the Advanced FTP V2 connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. Specify a Secure Agent.
Host	The host name or IP address of the FTP server.
Port	The port number to use for connecting to the FTP server. If left blank, the default port number 21 is used.
Username	User name to connect to the FTP server.
Password	Password to connect to the FTP server.
Folder Path	The directory to use after connecting to the FTP server.

Connection property	Description
Use passive mode	<p>Indicates whether the connection uses Passive or Active mode. Specify Yes to use Passive mode. Specify No to use Active mode.</p> <p>The default value is Yes.</p> <p>In Passive mode, the server does not need to connect back to a port on the connection client, which is a firewall-friendly mode. If you have problems with connecting to the server, you might want to change the mode to Passive by selecting Yes for this option. In Passive mode, depending on the FTP server, the connection may require high port range based on the port availability to transfer data.</p> <p>In Active mode, the server attempts to connect back to a port on the connection client to perform the data transfer.</p>
Data Connection Start Port	The starting port number to use for the data connection.
Data Connection End Port	The ending port number to use for the data connection.
Timeout	The number of seconds to wait when attempting to connect to the server. A timeout occurs if the connection cannot be established in the specified amount of time. If left blank, the default value of 120 seconds is used.
Connection Retry Attempts	The number of times to connect to retry the FTP connection if a connection cannot be established. This setting is used for both the initial connection and any reconnect attempts due to lost connections. If left blank, no retries will be attempted.
Connection Retry Interval	<p>The number of seconds to wait between each connection retry attempt.</p> <p>Note: For instance, if you want to retry to connect up to 10 times with a five second delay between retries, then specify 10 for the Connection Retry Attempts and 5 for the Connection Retry Interval.</p>
Control Encoding	If left blank, the connection uses the ISO standard ISO-8859-1. If supported by the server, other encodings such as UTF-8 can be specified to support international characters.
List Parser	The list parser to use for the server connection. If the field is left blank, the Advanced FTP V2 Connector attempts to use the MLSD parser. If the MLSD parser is not supported by the server, the UNIX parser is used. If you experience problems listing directories, select a different list parser.
Date Format	This date format is applied if the server returns a date that is different from the selected list parser default. If your location requires a different date format (for example, d MMM yyyy), specify the date format in this field. Not all list parsers support the date format setting. List parsers that do not support the date format setting ignores any user specified value.
Recent Date Format	Specify the date format to use when parsing the recent last modified date for each file. The recent date format applies in UNIX-based systems and appears on entries less than a year old. If your location requires a specific date format (for example, d MMM HH:mm), specify that pattern in this field. Not all list parsers support the recent date format setting. List parsers that do not support the recent date format setting ignores any user-specified value.

Connection property	Description
Bandwidth	Controls the maximum amount of network resources used for file transfers. The value is applicable for file uploads and downloads. Default is 0. 0 indicates that the bandwidth is not restricted.
Bandwidth Unit	The unit of the network bandwidth used for file transfer. You can choose one of the following units: <ul style="list-style-type: none"> - Kilobytes per second (KBps) - Megabytes per second (MBps)

Note: Advanced FTP V2 connector doesn't support NTLM proxy authentication.

Advanced FTPS V2 connection properties

When you set up an Advanced FTPS V2 connection, you must configure the connection properties.

The following table describes the Advanced FTPS V2 connection properties:

Connection property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~ ` ! \$ % ^ & * () - + = { [] } \ ; , ' < , > . ? /
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Type	Select the Advanced FTPS V2 connection type.
Runtime Environment	The name of the runtime environment that contains the Secure Agent that you want to run the tasks.
Host	The host name or IP address of the server.
Port	The port number to use for connecting to the server. If left blank, the default port number is 21.
Username	User name to connect to the FTPS server.
Password	Password to connect to the FTPS server.
Folder Path	The directory to use after connecting to the server.

Connection property	Description
Use passive mode	<p>Indicates whether the connection uses Passive or Active mode. Specify Yes to use Passive mode. Specify No to use Active mode.</p> <p>The default value is Yes.</p> <p>In Passive mode, the server does not need to connect back to a port on the connection client, which is a firewall-friendly mode. If you have problems with connecting to the server, you might want to change the mode to Passive by selecting Yes for this option. In Passive mode, depending on the FTPS server, the connection may require high port range based on the port availability to transfer data.</p> <p>In Active mode, the server attempts to connect back to a port on the connection client to perform the data transfer.</p>
Data Connection Start Port	The starting port number to use for the data connection.
Data Connection End Port	The ending port number to use for the data connection.
Timeout	The number of seconds to wait when attempting to connect to the server. A timeout occurs if the connection cannot be established in the specified amount of time. If left blank, the default value of 120 seconds is used.
Connection Retry Attempts	The number of times to connect to retry the Advanced FTP V2 connection if a connection cannot be established. This setting is used for both the initial connection and any reconnect attempts due to lost connections. If left blank, no retries will be attempted.
Connection Retry Interval	<p>The number of seconds to wait between each connection retry attempt.</p> <p>Note: For instance, if you want to retry to connect up to 10 times with a five second delay between retries, then specify 10 for the Connection Retry Attempts and 5 for the Connection Retry Interval.</p>
Control Encoding	If left blank, the connection uses the ISO standard ISO-8859-1. If supported by the server, other encodings like UTF-8 can be specified to support international characters.
Trusted Server	Specify whether the FTPS server is a trusted server. The Advanced FTP V2 Connector only supports a trusted server.
List Parser	The list parser to use for the server connection. If the field is empty, the Advanced FTP V2 Connector tries to use the MLSD parser. If the server does not support the MLSD parser, the connector uses the UNIX parser. If you experience problems listing directories, select a different list parser.
Date Format	This date format is applied if the server returns a date that is different from the selected list parser default. If your location requires a different date format (for example, d MMM yyyy), specify the date format in this field. Not all list parsers support the date format setting. List parsers that do not support the date format setting ignores any user specified values.
Recent Date Format	Specify the date format to use when parsing the recent last modified date for each file. The recent date format applies in UNIX-based systems and appears on entries less than a year old. If your location requires a specific date format (for example, d MMM HH:mm), specify that pattern in this field. Not all list parsers support the recent date format setting. List parsers that do not support the recent date format setting ignores any user-specified values.

Connection property	Description
Connection Type	Indicates if the connection type is IMPLICIT_SSL or EXPLICIT_SSL. <ul style="list-style-type: none"> - IMPLICIT_SSL. The connection automatically starts as an SSL connection. - EXPLICIT_SSL. After initial authentication with the FTPS server, the connection is encrypted with SSL or TLS depending on the security protocol you select. Default is IMPLICIT_SSL.
SecurityProtocol	Indicates whether SSL or TLS is used for EXPLICIT_SSL connections. Default is SSL.
Key Store File	The path and file name of the keystore file. The keystore file contains the certificates to authenticate the FTPS server.
Key Store Password	The password for the keystore file required to access the Trusted Server Certificate Store.
Key Alias	The alias of the individual key.
Key Store Type	Indicates if the type of the keystore is Java KeyStore (JKS) or Public Key Cryptology Standard (PKCS12). Default is JKS.
Bandwidth	Controls the maximum amount of network resources used for file transfers. The value is applicable for file uploads and downloads. Default is 0. 0 indicates that the bandwidth is not restricted.
Bandwidth Unit	The unit of the network bandwidth used for file transfer. You can choose one of the following units: <ul style="list-style-type: none"> - Kilobytes per second (KBps) - Megabytes per second (MBps)

Note: Advanced FTPS V2 connector doesn't support NTLM proxy authentication.

Advanced SFTP V2 connection properties

When you set up an Advanced SFTP V2 connection, you must configure the connection properties.

The following table describes the Advanced SFTP V2 connection properties:

Connection property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { }] \ : ; " ' < , > . ? /
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Type	Select the Advanced SFTP V2 connection type.

Connection property	Description
Runtime Environment	The name of the runtime environment that contains the Secure Agent that you want to run the tasks.
Host	The host name or IP address of the server. The host name is case insensitive and must be unique within the domain. The name cannot exceed 24 characters. It can contain letters (A to Z), digits (0 to 9), period (.) special character, and minus (-) sign.
Port	The port number to use to connect to the server. Default is 21.
Username	User name to connect to the SFTP server.
Password	Password to connect to the SFTP server.
Folder Path	The directory to use after connecting to the server.
Timeout	The number of seconds to wait when attempting to connect to the server. A timeout occurs if the connection cannot be established in the specified amount of time. If left blank, the default value of 120 seconds is used.
Connection Retry Attempts	The number of times to connect to retry the SFTP connection if a connection cannot be established. This setting is used for both the initial connection and any reconnect attempts due to lost connections. If left blank, no retries will be attempted.
Connection Retry Interval	The number of seconds to wait between each connection retry attempt. For example, if you want to retry to connect up to 10 times with a five second delay between retries, then specify 10 for the Connection Retry Attempts and 5 for the Connection Retry Interval.
Private Key File	The name of the SSH private key file along with the path where the file is stored. Ensure that the file path is on the machine that hosts the Secure Agent. For example, C:/SSH/my_keys/key.ppk
Private Key Passphrase	Specify the passphrase to encrypt the SSH private key.
Use Curve Kex Algorithm	Enable additional key exchange algorithms such as curve, and keyed-hash algorithm such as, -hmac-sha2-512, and -hmac-sha2-256.
Bandwidth	Controls the maximum amount of network resources used for file transfers. The value is applicable for file uploads and downloads. Default is 0. 0 indicates that the bandwidth is not restricted.
Bandwidth Unit	The unit of the network bandwidth used for file transfer. You can choose one of the following units: <ul style="list-style-type: none"> - Kilobytes per second (KBps) - Megabytes per second (MBps)
Use File Integration Proxy Server	The connector connects to the SFTP server through the file integration proxy server. Verify that the following prerequisites are met: <ul style="list-style-type: none"> - You must have the File Integration Service license to use this option. - You must define a proxy server in File Servers. - If you don't have the File Integration Service proxy, you need to use the agent proxy through the proxy.ini file.

Connection property	Description
Proxy Server Host Name	Host name or IP address of the outgoing File Integration Service proxy server.
Proxy Server Port	Port number of the outgoing File Integration Service proxy server.

Note: Advanced SFTP V2 connector doesn't support NTLM proxy authentication.

Amazon Kinesis connection properties

The Amazon Kinesis connection is a messaging connection. Use the Amazon Kinesis connection to access Amazon Kinesis Data Streams or Amazon Kinesis Data Firehose as targets.

Amazon Kinesis Firehose connection properties

When you set up an Amazon Kinesis Firehose connection, you must configure the connection properties.

The following table describes the Amazon Kinesis Firehose connection properties:

Property	Description
Connection Name	Name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
Description	Optional. Description that you can use to identify the connection. The description cannot exceed 4,000 characters.
Type	The Amazon Kinesis connection type. If you do not see the Amazon Kinesis connection type, go to the Add-On Connectors page to enable the connector.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Service	The type of Kinesis Service that you want to use. Select Kinesis Firehose .
AWS Access Key ID	The access key ID of the Amazon AWS user account.
AWS Secret Access Key	The secret access key for the Amazon AWS user account.

Property	Description
Region	<p>Region where the endpoint for your service is available. You can select one of the following values:</p> <ul style="list-style-type: none"> - us-east-2. Indicates the US East (Ohio) region. - us-east-1. Indicates the US East (N. Virginia) region. - us-west-1. Indicates the US West (N. California) region. - us-west-2. Indicates the US West (Oregon) region. - ap-northeast-1. Indicates the Asia Pacific (Tokyo) region. - ap-northeast-2. Indicates the Asia Pacific (Seoul) region. - ap-northeast-3. Indicates the Asia Pacific (Osaka-Local) region. - ap-south-1. Indicates the Asia Pacific (Mumbai) region. - ap-southeast-1. Indicates the Asia Pacific (Singapore) region. - ap-southeast-2. Indicates the Asia Pacific (Sydney) region. - ca-central-1. Indicates the Canada (Central) region. - cn-north-1. Indicates the China (Beijing) region. - cn-northwest-1. Indicates the China (Ningxia) region. - eu-central-1. Indicates the EU (Frankfurt) region. - eu-west-1. Indicates the EU (Ireland) region. - eu-west-2. Indicates the EU (London) region. - eu-west-3. Indicates the EU (Paris) region. - sa-east-1. Indicates the South America (São Paulo) region. - us-gov-west-1. Indicates AWS GovCloud (US-West) region. - us-gov-east-1. Indicates AWS GovCloud (US-East) region. <p>A streaming ingestion and replication task does not support ap-northeast-3 region.</p>
Connection Timeout (ms)	<p>Optional. Number of milliseconds that the Data Ingestion and Replication service waits to establish a connection to the Kinesis Firehose after which it times out.</p> <p>Default is 10,000 milliseconds.</p>
AWS Credential Profile Name	<p>An AWS credential profile defined in the credentials file.</p> <p>A mapping accesses the AWS credentials through the profile name at run time. If you do not provide an AWS credential profile name, the mapping uses the access key ID and secret access key that you specify when you create the connection.</p>
ARN of IAM Role	<p>The Amazon Resource Name specifying the role of an IAM user. Applies to Cross-Account IAM Roles authentication.</p>
External ID	<p>The external ID for an IAM role is an additional restriction that you can use in an IAM role trust policy to designate who can assume the IAM role. Applies to Cross-Account IAM Roles authentication.</p>

Amazon Kinesis Streams connection properties

When you set up an Amazon Kinesis Streams connection, you must configure the connection properties.

The following table describes the Amazon Kinesis Streams connection properties:

Property	Description
Connection Name	<p>Name of the connection.</p> <p>The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:</p> <p>~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /</p>
Description	<p>Optional. Description that you can use to identity the connection.</p> <p>The description cannot exceed 4,000 characters.</p>
Type	<p>The Amazon Kinesis connection type.</p> <p>If you do not see the Amazon Kinesis connection type, go to the Add-On Connectors page to install the connector.</p>
Runtime Environment	<p>Name of the runtime environment where you want to run the tasks.</p>
Service	<p>The type of Kinesis Service that you want to use. Select Kinesis Streams.</p>
AWS Access Key ID	<p>The access key ID of the Amazon AWS user account.</p>
AWS Secret Access Key	<p>The secret access key for your Amazon AWS user account.</p>
Region	<p>Region where the endpoint for your service is available. You can select one of the following values:</p> <ul style="list-style-type: none">- us-east-2. Indicates the US East (Ohio) region.- us-east-1. Indicates the US East (N. Virginia) region.- us-west-1. Indicates the US West (N. California) region.- us-west-2. Indicates the US West (Oregon) region.- ap-northeast-1. Indicates the Asia Pacific (Tokyo) region.- ap-northeast-2. Indicates the Asia Pacific (Seoul) region.- ap-northeast-3. Indicates the Asia Pacific (Osaka-Local) region.- ap-south-1. Indicates the Asia Pacific (Mumbai) region.- ap-southeast-1. Indicates the Asia Pacific (Singapore) region.- ap-southeast-2. Indicates the Asia Pacific (Sydney) region.- ca-central-1. Indicates the Canada (Central) region.- cn-north-1. Indicates the China (Beijing) region.- cn-northwest-1. Indicates the China (Ningxia) region.- eu-central-1. Indicates the EU (Frankfurt) region.- eu-west-1. Indicates the EU (Ireland) region.- eu-west-2. Indicates the EU (London) region.- eu-west-3. Indicates the EU (Paris) region.- sa-east-1. Indicates the South America (São Paulo) region.- us-gov-west-1. Indicates AWS GovCloud (US-West) region.- us-gov-east-1. Indicates AWS GovCloud (US-East) region. <p>A streaming ingestion and replication task does not support ap-northeast-3 region.</p>

Property	Description
Connection TimeOut (ms)	Optional. Number of milliseconds that the Data Ingestion and Replication service waits to establish a connection to the Kinesis Streams after which it times out. Default is 10,000 milliseconds.
AWS Credential Profile Name	An AWS credential profile defined in the credentials file. A mapping accesses the AWS credentials through the profile name at run time. If you do not provide an AWS credential profile name, the mapping uses the access key ID and secret access key that you specify when you create the connection.
ARN of IAM Role	The Amazon Resource Name specifying the role of an IAM user. Applies to Cross-Account IAM Roles authentication.
External ID	The external ID for an IAM role is an additional restriction that you can use in an IAM role trust policy to designate who can assume the IAM role. Applies to Cross-Account IAM Roles authentication.

AWS Credential Profile

You can define AWS credential profiles in the credentials file. Each credential profile consists of secret access key and access key ID.

Users can use the AWS credential profile names to use different AWS credentials at run time than the AWS credentials that they specify when they create an Amazon Kinesis connection with an Amazon Kinesis Streams as a source and target and Amazon Kinesis Firehose as a target.

Create AWS credentials for the users, such as access key ID and secret access key. Users can select an authentication type while creating an Amazon Kinesis connection, such as AWS credential profile. The default authentication type is AWS credential profile.

Generate an Access Key ID and Secret Access Key for the users in AWS.

Amazon Redshift V2 connection properties

Create an Amazon Redshift V2 connection to read from or write data to Amazon Redshift.

Prepare for authentication

You can configure **Default** and **Redshift IAM Authentication via AssumeRole** authentication types in an Amazon Redshift V2 connection to connect to Amazon Redshift. Additionally, you need to complete the S3 staging prerequisites to access S3 resources. You can also configure encryption, if required, to connect to Amazon Redshift.

Note: Application ingestion and replication and database ingestion and replication tasks do not support Redshift IAM authentication via AssumeRole unless you use an EC2 instance to assume the role.

See the following sections for a summary of the authentication, staging, and encryption prerequisites.

Authentication prerequisites

Before you begin, you need to have a registered user account with Amazon Redshift.

Get the minimum required details from your Amazon Redshift account from the AWS Console for the authentication type that you want to configure, as listed in the following table:

Default authentication	Redshift IAM Authentication via Assume Role
<ul style="list-style-type: none"> - JDBC URL - User name - Password 	<ul style="list-style-type: none"> - JDBC URL - User name - Database name - Cluster identifier - Redshift IAM role ARN*
<p>*To use the Redshift IAM role ARN, configure the Redshift IAM role ARN with the required trust policies to generate temporary security credentials to access Amazon Redshift.</p> <p>For instructions, see "Configure an assume role for Amazon Redshift" on page 31.</p>	

Staging prerequisites

To enable staging on Amazon S3 and to gain access to S3 resources when you read or write data, you need to configure the staging properties in the Amazon Redshift V2 connection.

The following table summarizes the staging options that you can configure in the connection for both default and Redshift IAM Authentication via AssumeRole authentication and the tasks that you need to perform to get the required details for S3 staging:

S3 staging options	Tasks
Generate temporary credentials for the IAM user who assumes the S3 IAM role to access S3 staging.	<p>AWS configurations</p> <p>Enable IAM users to assume an S3 IAM role and generate temporary credentials.</p> <p>For instructions, see the following references:</p> <ul style="list-style-type: none"> - "Generate temporary security credentials using AssumeRole for Amazon S3 staging" on page 35. - Using an assume role for Amazon S3 resources How-To Library article. <p>Redshift V2 connection configurations</p> <ul style="list-style-type: none"> - Enter the value of the S3 IAM Role ARN. - Enter the S3 Access Key ID and S3 Secret Access Key values.
Generate temporary security credentials for an EC2 instance that assumes an S3 IAM role to access S3 staging.	<p>AWS configurations</p> <p>Define an EC2 instance to assume an S3 IAM role and generate the temporary credentials for S3 staging.</p> <p>For instructions, see "Generate temporary security credentials using AssumeRole for EC2" on page 37.</p> <p>Redshift V2 connection configurations</p> <p>Configure the following minimum required properties:</p> <ul style="list-style-type: none"> - Enable Use EC2 Role to Assume Role. - Enter the value of the S3 IAM Role ARN.

S3 staging options	Tasks
Generate the S3 access and secret access keys for the IAM user with access to the S3 bucket.	<p>AWS configurations To generate the credentials, perform the following tasks:</p> <ol style="list-style-type: none"> 1. "Create a minimal Amazon IAM policy" on page 30. 2. Create an IAM user, assign the policy to that user, and then generate the S3 access key ID and S3 secret access key in the AWS console. <p>For more information about how to create an IAM user and generate keys, see the AWS documentation.</p> <p>Redshift V2 connection configurations Enter the S3 Access Key ID and S3 Secret Access Key values.</p>
Configure IAM authentication	<p>AWS configurations If you have an EC2 instance, and do not want to specify the keys or use the IAM role ARN, then assign the minimum policy to the EC2 with access to the S3 bucket.</p> <p>For instructions, see "Configure IAM authentication" on page 31.</p> <p>Redshift V2 connection configurations In this case, you do not need to enable or specify any of the staging properties in the connection.</p>

Encryption prerequisites

To configure client-side and server-side encryption for the Default authentication and Redshift IAM authentication via AssumeRole during staging, see ["Enable encryption" on page 37.](#)

Create a minimal Amazon IAM policy

To stage the data in Amazon S3, you need to create an IAM policy with the minimum required permissions to access the S3 resources.

You can either attach the policy to the IAM user and generate the S3 access key ID and S3 secret access keys to access S3 resources. Or, if you have an EC2 instance, you can assign the minimum policy to the EC2 instance to access the S3 bucket for staging.

You need the following minimum required permissions in the policy:

- PutObject
- GetObject
- DeleteObject
- ListBucket

You can use the following sample Amazon IAM policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:DeleteObject",
        "s3:ListBucket",
        "s3:ListBucketMultipartUploads"
      ],
      "Resource": [
        "arn:aws:s3:::<bucket_name>/*",

```

```

    "arn:aws:s3:::<bucket_name>"
  ]
}
]
}

```

Note: The **Test Connection** does not validate the IAM policy assigned to users. Hence, ensure that the policy assigned to the user is valid.

Configure IAM authentication

Configure AWS Identity and Access Management (IAM) authentication and create a minimal Amazon IAM policy for both the EC2 role and Redshift role.

For instructions, see the following How-to-Library article: [Configuring AWS IAM Authentication](#)

Configure an assume role for Amazon Redshift

To use the Redshift IAM role ARN, configure the Redshift IAM role ARN with the required trust policies to generate temporary security credentials to access Amazon Redshift.

You can use one of the following options to generate the temporary security credentials:

AWS configurations	Connection details
Option 1. Configure an AssumeRole to enable an IAM user.	To use the AssumeRole for the IAM user, specify the following IAM user details: <ul style="list-style-type: none"> - Redshift Access Key ID - Redshift Secret Access Key - Redshift IAM Role ARN
Option 2. Define an EC2 instance to assume a Redshift IAM role.	To use the AssumeRole for Amazon EC2: <ul style="list-style-type: none"> - Specify the Redshift IAM Role ARN value. - Enable the Use EC2 Role to Assume Role check box.

For application ingestion and replication tasks and database ingestion and replication tasks, use Option 2 to have an EC2 role assume the Redshift IAM role.

For more information about configuring an AssumeRole, see the following How-to-Library article: [Configure AssumeRole authentication for Amazon Redshift V2 Connector](#)

Generate the temporary security credentials based on your requirement.

Generate temporary security credential policies for Amazon Redshift

To use the temporary security credentials to connect to Amazon Redshift, both the IAM user and IAM role require policies.

The following section lists the policies required for the IAM user and IAM role:

IAM user

An IAM user must have the `sts:AssumeRole` policy to use the temporary security credentials in the same or different AWS account. The IAM user credentials are used to key-in the Redshift access key and Redshift secret key in the connection properties.

The following sample policy allows an IAM user to use the temporary security credentials in an AWS account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::<ACCOUNT-HYPHENS>:role/<REDSHIFT-IAM-ROLE-NAME>"
    }
  ]
}
```

Redshift IAM role trust policy

The Redshift IAM role policy pertains to the role that is specified in the Redshift IAM Role ARN. An IAM role must have a trust policy attached with it to allow the IAM user to access Redshift using the temporary security credentials.

The following policy is a sample trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": { "AWS": "arn:aws:iam::AWS-account-ID:<IAM-USER>" },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For example, you can specify the role or user in the following format:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<AWS-account>:role/<name-of-the-role>"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<AWS-account>:user/<name-of-the-user>"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Minimum permission policies of the Redshift IAM role

The following policy shows the permissions required to the Redshift IAM Role, which will be assumed by an IAM user to connect to the Redshift database using an existing Amazon Redshift user:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```



```

"Effect": "Allow",
"Action": [
    "redshift:GetClusterCredentials",
    "redshift:DescribeClusters"
],
"Resource": [
    "arn:aws:redshift:<REGION>:<ACCOUNT-ID>:dbuser:<Cluster_Identifier>/<USER_NAME>",
    "arn:aws:redshift:<REGION>:<ACCOUNT-ID>:dbname:<Cluster_Identifier>/<DATABASE_NAME>"
]
}
}
}
}

```

The following policy shows the permissions needed to be attached to the Redshift IAM Role, which will be assumed by an IAM user to connect to the Redshift database with a newly created user by the `Auto create DBUser` check box:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "redshift:GetClusterCredentials",
        "redshift:DescribeClusters",
        "redshift:CreateClusterUser",
        "redshift:JoinGroup"
      ],
      "Resource": [
        "arn:aws:redshift:<REGION>:<ACCOUNT-ID>:dbuser:<Cluster_Identifier>/<USER_NAME>",
        "arn:aws:redshift:<REGION>:<ACCOUNT-ID>:dbname:<Cluster_Identifier>/<DATABASE_NAME>",
        "arn:aws:redshift:<REGION>:<ACCOUNT-ID>:dbgroup:<Cluster_Identifier>/<GROUP_NAME>"
      ]
    }
  ]
}

```

Generate temporary security credentials using AssumeRole for EC2

You can use temporary security credentials using AssumeRole for an Amazon EC2 role to connect to Amazon Redshift from the same or different AWS accounts.

The Amazon EC2 role can assume another IAM role from the same or different AWS account without requiring a Redshift access key and Redshift secret key.

Consider the following prerequisites when you use temporary security credentials using AssumeRole for EC2:

- To use temporary security credentials using AssumeRole for EC2, install the Secure Agent on an AWS service such as Amazon EC2.
- The EC2 role attached to the AWS EC2 service must not have access to Amazon Redshift but needs to have permission to assume another IAM role.
- The IAM role that needs to be assumed by the EC2 role must have a permission policy and a trust policy attached to it.

To configure an EC2 role to assume the IAM Role provided in the Redshift IAM Role ARN connection property, select the **Use EC2 Role to Assume Role** check box in the connection properties.

EC2 service role trust policy

The following is a sample trust policy that is defined in a trust relationship of the EC2 role attached to the EC2 instance:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

The following is a sample trust policy of the Redshift IAM role when you enable EC2 assume role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "redshift.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<Account-ID:role>/ec2_role_attached_to_ec2_instance"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

The permission policy that is required to be attached to the EC2 instance is same as the policy defined for the IAM user.

Configure an assume role for Amazon S3 staging

To configure AssumeRole authentication for S3 staging, you need to attach the minimum permission policies and trust policies for the IAM user and IAM role in the AWS console.

An IAM user can use the AssumeRole to temporarily gain access to the Amazon S3 resources. For more information about using an assume role for Amazon S3 resources, you can also refer to the How-to-Library article: [Using an assume role for Amazon S3 resources](#)

You can generate temporary security credentials using AssumeRole for Amazon S3 staging to access the Amazon S3 staging bucket. If you want EC2 instances to assume an IAM role to gain access to the S3 staging bucket securely, use the temporary security credentials generated using AssumeRole for EC2 instances.

Note: Do not use the root user credentials of the AWS account to generate the temporary security credentials. You need to use the credentials of an IAM user to generate the temporary security credentials.

Generate the temporary security credentials based on your requirement.

Generate temporary security credentials using AssumeRole for Amazon S3 staging

You can use the temporary security credentials using AssumeRole to access the Amazon S3 staging bucket from the same or different AWS accounts.

Ensure that you have the **sts:AssumeRole** permission and a trust relationship established within the AWS accounts to use the temporary security credentials. The trust relationship is defined in the trust policy of the IAM role when you create the role. The IAM role adds the IAM user as a trusted entity allowing the IAM users to use the temporary security credentials and access the AWS accounts. For more information about how to establish the trust relationship, see the AWS documentation.

When the trusted IAM user requests for the temporary security credentials, the AWS Security Token Service (AWS STS) dynamically generates the temporary security credentials that are valid for a specified period and provides the credentials to the trusted IAM users. The temporary security credentials consist of access key ID, secret access key, and secret token.

To use the dynamically generated temporary security credentials, provide the value of the **S3 IAM Role ARN** connection property when you create an Amazon Redshift V2 connection. The IAM Role ARN uniquely identifies the AWS resources. Then, specify the time duration in seconds during which you can use the temporarily security credentials in the **Temporary Credential Duration** advanced source and target properties.

External ID

You can specify the external ID for a more secure access to the Amazon S3 bucket when the Amazon S3 bucket is in a different AWS account than the IAM user or EC2 instance.

Note: Application ingestion and replication and database ingestion and replication tasks do not support use of External ID.

You can optionally specify the external ID in the AssumeRole request to the AWS Security Token Service (STS).

The external ID must be a string. The following sample shows an external ID condition in the assumed IAM role's trust policy:

```
"Statement": [
  {
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::AWS_Account_ID : user/user_name"
    },
    "Action": "sts:AssumeRole",
    "Condition": {
      "StringEquals": {
        "sts:ExternalId": "dummy_external_id"
      }
    }
  }
]
```

```
}
]
```

Temporary security credentials policy

To use the temporary security credentials to access the Amazon S3 staging bucket, both the IAM user and IAM role require policies.

The following section lists the policies required for the IAM user and IAM role:

IAM user

An IAM user must have the `sts:AssumeRole` policy to use the temporary security credentials in the same or different AWS account.

The following sample policy allows an IAM user to use the temporary security credentials in an AWS account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::<ACCOUNT-HYPHENS>:role/<ROLE-NAME>"
    }
  ]
}
```

The following sample policy allows an IAM user for the China region to use the temporary security credentials in an AWS account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws-cn:iam::<ACCOUNT-HYPHENS>:role/<ROLE-NAME>"
    }
  ]
}
```

IAM role

An IAM role must have a `sts:AssumeRole` policy and a trust policy attached with the IAM role to allow the IAM user to access the Amazon S3 bucket using the temporary security credentials. The policy specifies the Amazon S3 bucket that the IAM user can access and the actions that the IAM user can perform. The trust policy specifies the IAM user from the AWS account that can access the Amazon S3 bucket.

The following policy is a sample trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::AWS-account-ID:<ROLE-NAME>"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Temporary security credentials for KMS

To use the temporary security credentials with AWS Key Management Service (AWS KMS)-managed customer master key and enable encryption with KMS, you must create a KMS policy.

You can perform the following operations to use the temporary security credentials and enable encryption with KMS:

- `GenerateDataKey`

- DescribeKey
- Encrypt
- Decrypt
- ReEncrypt

You can use the following sample policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:GenerateDataKey", "kms:DescribeKey", "kms:Encrypt", "kms:Decrypt",
        "kms:ReEncrypt*"
      ],
      "Resource": ["arn:aws:kms:region:account:key/<KMS_key>"]
    }
  ]
}
```

When you configure KMS and access an Amazon S3 endpoint in the China region, use the following sample policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:GenerateDataKey", "kms:DescribeKey", "kms:Encrypt", "kms:Decrypt",
        "kms:ReEncrypt*"
      ],
      "Resource": ["arn:aws-cn:kms:region:account:key/<KMS_key>"]
    }
  ]
}
```

Generate temporary security credentials using AssumeRole for EC2

You can use temporary security credentials using AssumeRole for an Amazon EC2 role to access the Amazon S3 staging bucket from the same or different AWS accounts.

The Amazon EC2 role can assume another IAM role from the same or different AWS account without requiring a permanent access key and secret key. The Amazon EC2 role can also assume another IAM role from a different region.

Consider the following prerequisites when you use temporary security credentials using AssumeRole for EC2:

- To use temporary security credentials using AssumeRole for EC2, install the Secure Agent on an AWS service such as Amazon EC2.
- The EC2 role attached to the AWS EC2 service must not have access to Amazon S3 but needs to have permission to assume another IAM role.
- The IAM role that needs to be assumed by the EC2 role must have a permission policy and a trust policy attached to it.

To configure an EC2 role to assume the IAM Role provided in the **IAM Role ARN** connection property, select the **Use EC2 Role to Assume Role** check box in the connection properties.

Enable encryption

You can enable client-side and server-side encryption in the Amazon Redshift V2 connection for staging data in Amazon S3.

Complete the prerequisites based on the type of encryption that you want to configure in the Amazon Redshift V2 connection.

Client-side encryption

Client-side encryption requires a 256-bit AES encryption key in the Base64 format. You can generate a key using a third-party tool.

Specify the key value in the **Master Symmetric Key** field when you create an Amazon Redshift V2 connection.

Server-side encryption

To enable server-side encryption, create an AWS Key Management Service (AWS KMS)-managed customer master key.

Generate the customer master key ID for the same region where your Amazon S3 staging bucket resides. For more information about generating a customer master key, see the AWS documentation.

To enable encryption with the customer master key, you need to create a minimal KMS policy. You can specify the customer master key ID when you create an Amazon Redshift V2 connection.

Note: You cannot configure server-side encryption with the master symmetric key and client-side encryption with the customer master key.

Create a minimal policy for using AWS KMS

To use the AWS Key Management Service (AWS KMS)-managed customer master key and enable the encryption with KMS, you must create a KMS policy.

You can perform the following operations to enable encryption with KMS:

- GenerateDataKey
- DescribeKey
- Encrypt
- Decrypt
- ReEncrypt

Sample policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:GenerateDataKey", "kms:DescribeKey", "kms:Encrypt", "kms:Decrypt",
        "kms:ReEncrypt*"
      ],
      "Resource": ["arn:aws:kms:region:account:key:<KMS_key>"]
    }
  ]
}
```

When you configure KMS and access an Amazon S3 endpoint in the China region, use the following sample policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:GenerateDataKey", "kms:DescribeKey", "kms:Encrypt", "kms:Decrypt",
        "kms:ReEncrypt*"
      ],
      "Resource": ["arn:aws-cn:kms:region:account:key:<KMS_key>"]
    }
  ]
}
```

Connect to Amazon Redshift

Let's configure the Amazon Redshift V2 connection properties to connect to Amazon Redshift.

Before you begin

Check out [“Prepare for authentication” on page 28](#) to learn about the authentication requirements before you configure a connection.

Connection details

The following table describes the basic connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - , Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	Amazon Redshift V2
Runtime Environment	Name of the runtime environment where you want to run tasks. You cannot run an application ingestion and replication task, database ingestion and replication task, file ingestion and replication task, or streaming ingestion and replication task on a Hosted Agent or serverless runtime environment.

Authentication types

You can configure default and Redshift IAM AssumeRole authentication types to access Amazon Redshift.

Note: Application ingestion and replication tasks and database ingestion and replication tasks do not support Redshift IAM AssumeRole authentication without an EC2 instance.

Select the required authentication method and then configure the authentication-specific parameters.

Default authentication

The following table describes the basic connection properties for default authentication:

Properties	Description
JDBC URL	The JDBC URL to connect to the Amazon Redshift cluster. You can get the JDBC URL from your Amazon AWS Redshift cluster configuration page. Enter the JDBC URL in the following format: <code>jdbc:redshift://<cluster_endpoint>:<port_number>/<database_name></code> , where the endpoint includes the Redshift cluster name and region. For example, <code>jdbc:redshift://infa-rs-cluster.abc.us-west-2.redshift.amazonaws.com:5439/rsdb</code> In the example, <ul style="list-style-type: none">- <code>infa-rs-qa-cluster</code> is the name of the Redshift cluster.- <code>us-west-2.redshift.amazonaws.com</code> is the Redshift cluster endpoint, which is the US West (Oregon) region.- <code>5439</code> is the port number for the Redshift cluster.- <code>rsdb</code> is the specific database instance in the Redshift cluster to which you want to connect.
Username	User name of your database instance in the Amazon Redshift cluster.

Properties	Description
Password	Password of the Amazon Redshift database user.
Use EC2 Role to Assume Role	<p>Enables the EC2 instance that assumes an S3 IAM role to access the S3 resources to stage data using the temporary security credentials.</p> <p>The EC2 role must have a policy attached with permissions to assume an S3 IAM role. The S3 IAM role and the EC2 instance can be in the same or different AWS account.</p> <p>Select the check box to enable the EC2 role to assume an S3 IAM role specified in the S3 IAM Role ARN option to access the S3 resources for staging data.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks. By default, this check box is not selected.</p> <p>For instructions, see "Generate temporary security credentials using AssumeRole for EC2" on page 37.</p>
S3 IAM Role ARN	<p>The Amazon Resource Number (ARN) of the IAM role assumed by the IAM user or EC2 to use the dynamically generated temporary security credentials to stage data in Amazon S3.</p> <p>This property applies when you want to generate temporary security credentials to access the S3 staging buckets by using either the EC2 instance or the IAM user who assumes the S3 IAM role.</p> <p>Specify the S3 IAM role name to use the temporary security credentials to access the Amazon S3 staging bucket.</p> <p>For more information about how to get the ARN of the S3 IAM role, see the AWS documentation.</p> <p>Note: If you use the connection for application ingestion and replication or database ingestion and replication tasks that use role-based authentication, but not the default role for the AWS cluster, specify an IAM role ARN. If you use the default role, leave this field blank.</p>

Advanced settings

The following table describes the advanced connection properties for default authentication:

Properties	Description
S3 Access Key ID	<p>Access key of the IAM user to access the Amazon S3 staging bucket.</p> <p>Enter the access key ID when you use the following methods for S3 staging:</p> <ul style="list-style-type: none"> - When the IAM user has access to S3 staging. - When the IAM user who assumes the S3 IAM role uses the temporary security credentials to access S3. <p>You do not need to enter the S3 access key ID if you use IAM authentication or the assume role for EC2 to access S3.</p> <p>Note: If you use the connection for application ingestion and replication or database ingestion and replication tasks that use key-based authentication, provide the access key value.</p>
S3 Secret Access Key	<p>Secret access key to access the Amazon S3 staging bucket.</p> <p>The secret key is associated with the access key and uniquely identifies the account.</p> <p>Enter the secret access key value when you use following methods for S3 staging:</p> <ul style="list-style-type: none"> - When the IAM user has access to S3 staging. - When the IAM user who assumes the S3 IAM role uses the temporary security credentials to access S3. <p>You do not need to enter the S3 secret access key if you use IAM authentication or the assume role for EC2 to access S3.</p> <p>Note: If you use the connection for application ingestion and replication or database ingestion and replication tasks that use key-based authentication, provide the access key value.</p>

Properties	Description
S3 VPC Endpoint Type	<p>The type of Amazon Virtual Private Cloud endpoint for Amazon S3.</p> <p>You can use a VPC endpoint to enable private communication with Amazon S3.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Default. Select if you do not want to use a VPC endpoint. - Interface Endpoint. Select to establish private communication with Amazon S3 through an interface endpoint which uses a private IP address from the IP address range of your subnet. It serves as an entry point for traffic destined to an AWS service.
Endpoint DNS Name for Amazon S3	<p>The DNS name for the Amazon S3 interface endpoint.</p> <p>Replace the asterisk symbol with the bucket keyword in the DNS name.</p> <p>Enter the DNS name in the following format:</p> <p>bucket.<DNS name of the interface endpoint></p> <p>For example, bucket.vpce-s3.us-west-2.vpce.amazonaws.com</p>
STS VPC Endpoint Type	<p>The type of Amazon Virtual Private Cloud endpoint for AWS Security Token Service.</p> <p>You can use a VPC endpoint to enable private communication with Amazon Security Token Service.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Default. Select if you do not want to use a VPC endpoint. - Interface Endpoint. Select to establish private communication with Amazon Security Token Service through an interface endpoint which uses a private IP address from the IP address range of your subnet.
Endpoint DNS Name for AWS STS	<p>The DNS name for the AWS STS interface endpoint.</p> <p>For example, vpce-01f22cc14558c241f-s8039x4c.sts.us-west-2.vpce.amazonaws.com</p>
KMS VPC Endpoint Type	<p>The type of Amazon Virtual Private Cloud endpoint for AWS Key Management Service.</p> <p>You can use a VPC endpoint to enable private communication with Amazon Key Management Service.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Default. Select if you do not want to use a VPC endpoint. - Interface Endpoint. Select to establish private communication with Amazon Key Management Service through an interface endpoint which uses a private IP address from the IP address range of your subnet.
Endpoint DNS Name for AWS KMS	<p>The DNS name for the AWS KMS interface endpoint.</p> <p>For example, vpce-0e722f5c721e19232-g2pkm2r7.kms.us-west-2.vpce.amazonaws.com</p>
External ID	<p>The external ID associated with the IAM role.</p> <p>You can specify the external ID if you want to provide a more secure access to the Amazon S3 bucket. The Amazon S3 staging bucket and the IAM role can be in the same or different AWS accounts.</p> <p>If required, you also have the option to specify the external ID in the AssumeRole request to the AWS Security Token Service (STS) using an external ID condition in the assumed IAM role's trust policy.</p> <p>For more information about using an external ID, see External ID when granting access to your AWS resources.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks.</p>

Properties	Description
Cluster Region	<p>The AWS cluster region in which the Redshift cluster resides.</p> <p>Select the cluster region from the list if you choose to provide a custom JDBC URL with a different cluster region from that specified in the JDBC URL field property. To continue to use the cluster region name specified in the JDBC URL field property, select None as the cluster region in this property.</p> <p>You can only read data from or write data to the cluster regions supported by the AWS SDK.</p> <p>Select one of the following cluster regions:</p> <p>None</p> <p>Asia Pacific(Mumbai)</p> <p>Asia Pacific(Seoul)</p> <p>Asia Pacific(Singapore)</p> <p>Asia Pacific(Sydney)</p> <p>Asia Pacific(Tokyo)</p> <p>Asia Pacific(Hong Kong)</p> <p>AWS GovCloud (US)</p> <p>AWS GovCloud (US-East)</p> <p>Canada(Central)</p> <p>China(Beijing)</p> <p>China(Ningxia)</p> <p>EU(Ireland)</p> <p>EU(Frankfurt)</p> <p>EU(Paris)</p> <p>EU(Stockholm)</p> <p>South America(Sao Paulo)</p> <p>Middle East(Bahrain)</p> <p>US East(N. Virginia)</p> <p>US East(Ohio)</p> <p>US West(N. California)</p> <p>US West(Oregon)</p> <p>Default is None.</p> <p>Note: A region value is required for application ingestion and replication tasks and database ingestion and replication tasks.</p>
Connection Environment SQL	<p>The SQL statement to set up the database environment that applies for the entire session.</p> <p>Separate multiple values with a semicolon (;).</p> <p>Specify only the configurations for the database environment in the SQL statement. Do not specify any DDL or DML commands in the SQL statement.</p>

Properties	Description
Master Symmetric Key	<p>A 256-bit AES encryption key in the Base64 format that enables client-side encryption to encrypt your data before you send them for staging in Amazon S3.</p> <p>For more information, see "Enable encryption" on page 37.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks.</p>
Customer Master Key ID	<p>The customer master key ID generated by AWS Key Management Service (AWS KMS) or the ARN of your custom key for cross-account access when you stage data in Amazon S3. The customer master key serves to encrypt your data at the destination before they are saved in Amazon S3. You can either enter the customer-generated customer master key ID or the default customer master key ID.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks.</p>

Redshift IAM Authentication via AssumeRole

The Redshift AssumeRole authentication enables the user to assume an IAM role or define an EC2 role configured with required trust policies to generate temporary security credentials to access Amazon Redshift.

Note: For application ingestion and replication tasks and database ingestion and replication tasks, you must use an EC2 role.

The following table describes the basic connection properties for Redshift IAM AssumeRole authentication:

Properties	Description
JDBC URL	<p>The JDBC URL to connect to the Amazon Redshift cluster.</p> <p>You can get the JDBC URL from your Amazon AWS Redshift cluster configuration page.</p> <p>Enter the JDBC URL in the following format:</p> <p><code>jdbc:redshift://<cluster_endpoint>:<port_number>/<database_name></code>, where the endpoint includes the Redshift cluster name and region.</p> <p>For example, <code>jdbc:redshift://infa-rs-cluster.abc.us-west-2.redshift.amazonaws.com:5439/rsdb</code></p> <p>In the example,</p> <ul style="list-style-type: none"> - <code>infa-rs-qa-cluster</code> is the name of the Redshift cluster. - <code>us-west-2.redshift.amazonaws.com</code> is the Redshift cluster endpoint, which is the US West (Oregon) region. - <code>5439</code> is the port number for the Redshift cluster. - <code>rsdb</code> is the specific database instance in the Redshift cluster to which you want to connect.
Username	User name of your database instance in the Amazon Redshift cluster.
Cluster Identifier	<p>The unique identifier of the cluster that hosts Amazon Redshift.</p> <p>Specify the Amazon Redshift cluster name.</p>
Database Name	Name of the Amazon Redshift database where the tables that you want to access are stored.

Properties	Description
Redshift IAM Role ARN	The Amazon Resource Number (ARN) of the IAM role assumed by EC2 to use the dynamically generated temporary security credentials to access Amazon Redshift. Enter the Redshift IAM role ARN to access the Amazon Redshift cluster.
Use EC2 Role to Assume Role	Enables the EC2 role to assume an IAM role, either to connect to Redshift or to stage data using the temporary security credentials: Connect to Redshift with IAM authentication using the EC2 role Select the check box to enable the EC2 role that assumes a Redshift IAM role specified in the Redshift IAM Role ARN field to access Amazon Redshift. The EC2 role must have a policy attached with permissions to assume a Redshift IAM role from the same or different account. Access S3 resources to stage data Select the check box to enable the EC2 role to assume an S3 IAM role specified in the S3 IAM Role ARN field and dynamically generate the temporary security credentials to access the S3 staging buckets. The EC2 role must have a policy attached with permissions to assume an S3 IAM role from the same or different AWS account.
S3 IAM Role ARN	The Amazon Resource Number (ARN) of the S3 IAM role assumed by the IAM user or EC2 to use the dynamically generated temporary security credentials to stage data in Amazon S3. This property applies when you want to generate the temporary security credentials to access the S3 staging buckets by using either the EC2 instance or the IAM user who assumes the S3 IAM role. Specify the S3 IAM role name to use the temporary security credentials to access the Amazon S3 staging bucket. For more information about how to get the ARN of the IAM role, see the AWS documentation . Note: If you use the connection for application ingestion and replication or database ingestion and replication tasks that uses role-based authentication, but not the default role for the AWS cluster, specify an IAM role ARN. If you use the default role, leave this field blank.

Advanced settings

The following table describes the advanced connection properties for Redshift IAM AssumeRole authentication:

Properties	Description
Redshift Access Key ID	The access key of the IAM user that has permissions to assume the Redshift IAM AssumeRole ARN. This property doesn't apply to Amazon Redshift AssumeRole authentication with EC2 role.
Redshift Secret Access Key	The secret access key of the IAM user that has permissions to assume the Redshift IAM AssumeRole ARN. This property doesn't apply to Amazon Redshift AssumeRole authentication with EC2 role.

Properties	Description
Database Group	<p>The name of the database group to which you want to add the database user when you select the Auto Create DBUser option in this connection property.</p> <p>The user that you add to this database group inherits the specified group privileges.</p> <p>If you do not specify a database group name, the user is added to the public group and inherits its associated privileges.</p> <p>You can also enter multiple database groups, separated by a comma, to add the user to each of the specified database groups.</p>
Expiration Time	<p>The time duration that the password for the Amazon Redshift database user expires.</p> <p>Specify a value between 900 seconds and 3600 seconds.</p> <p>Default is 900.</p>
Auto Create DBUser	<p>Select to create a new Amazon Redshift database user at run time.</p> <p>The agent adds the user you specified in the Username field to the database group. The added user assumes the privileges assigned to the database group.</p> <p>Default is disabled.</p>
S3 Access Key ID	<p>Access key of the IAM user to access the Amazon S3 staging bucket.</p> <p>Enter the access key ID when you use the following methods for S3 staging:</p> <ul style="list-style-type: none"> - When the IAM user has access to S3 staging. - When the IAM user who assumes the S3 IAM role uses the temporary security credentials to access S3. <p>You do not need to enter the S3 access key ID if you use IAM authentication or the assume role for EC2 to access S3.</p> <p>Note: If you use the connection for application ingestion and replication or database ingestion and replication tasks that use key-based authentication, provide the access key value.</p>
S3 Secret Access Key	<p>Secret access key to access the Amazon S3 staging bucket.</p> <p>The secret key is associated with the access key and uniquely identifies the account.</p> <p>Enter the secret access key value when you use following methods for S3 staging:</p> <ul style="list-style-type: none"> - When the IAM user has access to S3 staging. - When the IAM user who assumes the S3 IAM role uses the temporary security credentials to access S3. <p>You do not need to enter the S3 secret access key if you use IAM authentication or the assume role for EC2 to access S3.</p> <p>Note: If you use the connection for application ingestion and replication or database ingestion and replication tasks that use key-based authentication, provide the access key value.</p>
S3 VPC Endpoint Type	<p>The type of Amazon Virtual Private Cloud endpoint for Amazon S3.</p> <p>You can use a VPC endpoint to enable private communication with Amazon S3.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Default. Select if you do not want to use a VPC endpoint. - Interface Endpoint. Select to establish private communication with Amazon S3 through an interface endpoint which uses a private IP address from the IP address range of your subnet. It serves as an entry point for traffic destined to an AWS service.
Endpoint DNS Name for Amazon S3	<p>The DNS name for the Amazon S3 interface endpoint.</p> <p>Replace the asterisk symbol with the bucket keyword in the DNS name.</p> <p>Enter the DNS name in the following format:</p> <p><code>bucket.<DNS name of the interface endpoint></code></p> <p>For example, <code>bucket.vpce-s3.us-west-2.vpce.amazonaws.com</code></p>

Properties	Description
STS VPC Endpoint Type	<p>The type of Amazon Virtual Private Cloud endpoint for AWS Security Token Service. You can use a VPC endpoint to enable private communication with Amazon Security Token Service.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Default. Select if you do not want to use a VPC endpoint. - Interface Endpoint. Select to establish private communication with Amazon Security Token Service through an interface endpoint which uses a private IP address from the IP address range of your subnet.
Endpoint DNS Name for AWS STS	<p>The DNS name for the AWS STS interface endpoint.</p> <p>For example, <code>vpce-01f22cc14558c241f-s8039x4c.sts.us-west-2.vpce.amazonaws.com</code></p>
KMS VPC Endpoint Type	<p>The type of Amazon Virtual Private Cloud endpoint for AWS Key Management Service. You can use a VPC endpoint to enable private communication with Amazon Key Management Service.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Default. Select if you do not want to use a VPC endpoint. - Interface Endpoint. Select to establish private communication with Amazon Key Management Service through an interface endpoint which uses a private IP address from the IP address range of your subnet.
Endpoint DNS Name for AWS KMS	<p>The DNS name for the AWS KMS interface endpoint.</p> <p>For example, <code>vpce-0e722f5c721e19232-g2pkm2r7.kms.us-west-2.vpce.amazonaws.com</code></p>
External ID	<p>The external ID associated with the IAM role.</p> <p>You can specify the external ID if you want to provide a more secure access to the Amazon S3 bucket when the Amazon S3 staging bucket is in same or different AWS accounts.</p> <p>If required, you also have the option to specify the external ID in the AssumeRole request to the AWS Security Token Service (STS) using an external ID condition in the assumed IAM role's trust policy.</p> <p>For more information about using an external ID, see External ID when granting access to your AWS resources.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks.</p>

Properties	Description
Cluster Region	<p>The AWS geographical region in which the Redshift cluster resides.</p> <p>Select the cluster region from the list if you choose to provide a custom JDBC URL with a different cluster region from that specified in the JDBC URL field property. To continue to use the cluster region name specified in the JDBC URL field property, select None as the cluster region in this property.</p> <p>You can only read data from or write data to the cluster regions supported by the AWS SDK.</p> <p>Select one of the following cluster regions:</p> <p>None</p> <p>Asia Pacific(Mumbai)</p> <p>Asia Pacific(Seoul)</p> <p>Asia Pacific(Singapore)</p> <p>Asia Pacific(Sydney)</p> <p>Asia Pacific(Tokyo)</p> <p>Asia Pacific(Hong Kong)</p> <p>AWS GovCloud (US)</p> <p>AWS GovCloud (US-East)</p> <p>Canada(Central)</p> <p>China(Beijing)</p> <p>China(Ningxia)</p> <p>EU(Ireland)</p> <p>EU(Frankfurt)</p> <p>EU(Paris)</p> <p>EU(Stockholm)</p> <p>South America(Sao Paulo)</p> <p>Middle East(Bahrain)</p> <p>US East(N. Virginia)</p> <p>US East(Ohio)</p> <p>US West(N. California)</p> <p>US West(Oregon)</p> <p>Default is None.</p> <p>Note: A region value is required for application ingestion and replication tasks and database ingestion and replication tasks.</p>
Connection Environment SQL	<p>The SQL statement to set up the database environment that applies for the entire session.</p> <p>Separate multiple values with a semicolon (;).</p> <p>Specify only the configurations for the database environment in the SQL statement. Do not specify any DDL or DML commands in the SQL statement.</p>

Properties	Description
Master Symmetric Key	<p>A 256-bit AES encryption key in the Base64 format that enables client-side encryption to encrypt your data before you send them for staging in Amazon S3.</p> <p>For more information, see "Enable encryption" on page 37.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks.</p>
Customer Master Key ID	<p>The customer master key ID generated by AWS Key Management Service (AWS KMS) or the ARN of your custom key for cross-account access when you stage data in Amazon S3. The customer master key serves to encrypt your data at the destination before they are saved in Amazon S3. You can either enter the customer-generated customer master key ID or the default customer master key ID.</p> <p>For more information about how to configure server-side encryption, see "Enable encryption" on page 37.</p> <p>This property doesn't apply to application ingestion and replication tasks and database ingestion and replication tasks.</p>

Proxy server settings

If your organization uses an outgoing proxy server to connect to the Internet, the Secure Agent connects to Informatica Intelligent Cloud Services through the proxy server.

You can configure the Secure Agent to use the proxy server on Windows and Linux. You can use only an unauthenticated proxy server.

To configure the proxy settings for the Secure Agent, use one of the following methods:

- Configure the Secure Agent through the Secure Agent Manager on Windows or shell command on Linux. For instructions, see "Configure the proxy settings on Windows" or "Configure the proxy settings on Linux" in *Getting Started* in the Data Integration help.
- Configure the JVM options for the DTM in the Secure Agent properties. For instructions, see the [Proxy server settings](#) Knowledge Base article.

Note: If you enable both HTTP and SOCKS proxies, SOCKS proxy is used by default. If you want to use HTTP proxy instead of SOCKS proxy, set the value of the **DisableSocksProxy** property to true in the System property.

Private communication with Amazon Redshift

If you do not want to expose your traffic to the public internet, you can enable private communication with Amazon Redshift by configuring a gateway endpoint on the AWS console.

To establish a private connection with Amazon Redshift, ensure that the Secure Agent is a part of the subnet in the AWS Virtual Private Cloud (VPC). You can create a gateway endpoint and stage the Amazon S3 data to Amazon Redshift.

To configure private communication to connect to Amazon Redshift, you need to perform the following tasks:

- Create a cluster subnet group.
- Create a Redshift-managed VPC endpoint.
- Configure the gateway endpoint.

You can then specify the gateway endpoint in the Amazon Redshift V2 connection properties.

For more information, see
[Configuring private communication with Amazon Redshift using the Amazon Redshift V2 Connector.](#)

Amazon S3 V2 connection properties

When you set up an Amazon S3 V2 connection, configure the connection properties.

The following table describes the Amazon S3 V2 connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Amazon S3 V2 connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks. You cannot run an application ingestion and replication task or a database ingestion and replication task on a Hosted Agent or serverless runtime environment.
Access Key	Access key to access the Amazon S3 bucket. Enter the access key value based on the following authentication methods: <ul style="list-style-type: none">- Basic authentication. Enter the actual access key value.- IAM authentication. Don't enter the access key value.- Temporary security credentials using assume role. Enter the secret access key of an IAM user with no permissions to access Amazon S3 bucket.- Assume role for EC2. Don't enter the access key value.- Credential profile file authentication. Don't enter the access key value.- Federated user single sign-on. Don't enter the secret access key value.
Secret Key	Secret access key to access the Amazon S3 bucket. The secret key is associated with the access key and uniquely identifies the account. Enter the secret access key value based on the following authentication methods: <ul style="list-style-type: none">- Basic authentication. Enter the actual access secret value.- IAM authentication. Don't enter the access secret value.- Temporary security credentials using assume role. Enter access secret of an IAM user with no permissions to access Amazon S3 bucket.- Assume role for EC2. Don't enter the access key value.- Credential profile file authentication. Don't enter the access secret value.- Federated user single sign-on. Don't enter the access secret value.

Property	Description
IAM Role ARN	<p>The Amazon Resource Name (ARN) of the AWS Identity and Access Management (IAM) role assumed by the user to use the dynamically generated temporary security credentials.</p> <p>Enter the value of this property if you want to use the temporary security credentials to access the AWS resources.</p> <p>This property is not applicable to an application ingestion and replication task.</p> <p>Note: Even if you remove the IAM role that enables the agent to access the Amazon S3 bucket and create a connection, the test connection is successful.</p> <p>For more information about how to get the ARN of the IAM role, see the AWS documentation.</p>
External Id	Provides a more secure access to the Amazon S3 bucket when the Amazon S3 bucket is in a different AWS account.
Use EC2 Role to Assume Role	<p>Enables the EC2 role to assume another IAM role specified in the IAM Role ARN option.</p> <p>Note: The EC2 role must have a policy attached with a permission to assume an IAM role from the same or different account.</p> <p>By default, the Use EC2 Role to Assume Role check box is not selected.</p> <p>Note: Enter a value for the IAM Role ARN property when you enable this property for a streaming ingestion and replication task.</p>
Folder Path	<p>Bucket name or complete folder path to the Amazon S3 objects.</p> <p>For tasks other than application ingestion and replication and database ingestion and replication tasks, don't use a slash at the end of the folder path. For example, <code><bucket name>/<my folder name></code>.</p> <p>For application ingestion and replication and database ingestion and replication tasks, add a trailing slash. For example: <code><bucket name>/<my folder name>/</code>.</p>
Master Symmetric Key	<p>A 256-bit AES encryption key in the Base64 format when you use client-side encryption. You can generate a key using a third-party tool.</p> <p>Doesn't apply to an application ingestion and replication task, database ingestion and replication task, or streaming ingestion and replication task.</p>
Customer Master Key ID	<p>The customer master key ID or alias name generated by AWS Key Management Service (AWS KMS) or the Amazon Resource Name (ARN) of your custom key for cross-account access.</p> <p>You must generate the customer master key for the same region where the Amazon S3 bucket resides.</p> <p>You can specify the following master keys:</p> <ul style="list-style-type: none"> - Customer generated customer master key. Enables client-side or server-side encryption. - Default customer master key. Enables client-side or server-side encryption. Only the administrator user of the account can use the default customer master key ID to enable client-side encryption. <p>Doesn't apply to an application ingestion and replication task, database ingestion and replication task, or streaming ingestion and replication task.</p>
S3 Account Type	<p>The type of the Amazon S3 account.</p> <p>Select from the following options:</p> <ul style="list-style-type: none"> - Amazon S3 Storage. Enables you to use the Amazon S3 services. - S3 Compatible Storage. Enables you to use the endpoint for a third-party storage provider such as Scalify RING or MinIO. <p>Default is Amazon S3 storage.</p>

Property	Description
REST Endpoint	<p>The S3 storage endpoint required for S3 compatible storage. Enter the S3 storage endpoint in HTTP or HTTPS format. For example, http://s3.isv.scality.com.</p>
Region Name	<p>The AWS region of the bucket that you want to access. Select one of the following regions:</p> <ul style="list-style-type: none"> - Africa(Cape Town) - Asia Pacific(Mumbai) - Asia Pacific(Jakarta) - Asia Pacific (Osaka) - Asia Pacific(Seoul) - Asia Pacific(Singapore) - Asia Pacific(Sydney) - Asia Pacific(Tokyo) - Asia Pacific(Hong Kong) - AWS GovCloud(US) - AWS GovCloud(US-East) - Canada(Central) - China(Beijing) - China(Ningxia) - EU(Ireland) - EU(Frankfurt) - EU(London) - EU(Milan) - EU(Paris) - EU(Stockholm) - South America(Sao Paulo) - Middle East(Bahrain) - Middle East(UAE) - US East(N. Virginia) - US East(Ohio) - US ISO East - US ISOB East(Ohio) - US ISO West - US West(N. California) - US West(Oregon) <p>Default is US East(N. Virginia).</p>
Federated SSO IdP	<p>SAML 2.0-enabled identity provider for the federated user single sign-on to use with the AWS account.</p> <p>Amazon S3 V2 connector supports only the ADFS 3.0 identity provider. Select None if you don't want to use federated user single sign-on.</p> <p>Note: Federated user single sign-on is not applicable to application ingestion and replication tasks, database ingestion and replication tasks, and streaming ingestion and replication tasks.</p>
Other Authentication Type	<p>Select one the following authentication types:</p> <ul style="list-style-type: none"> - NONE - Credential Profile File Authentication <p>Select the Credential Profile File Authentication option to access the Amazon S3 credentials from a credential file that contains the access key and secret key.</p> <p>Enter the credential profile file path and the profile name to establish the connection with Amazon S3.</p> <p>You can use permanent IAM credentials or temporary session tokens when you configure the Credential Profile File Authentication.</p> <p>Default is NONE.</p>

Property	Description
Credential Profile File Path	<p>Specifies the credential profile file path.</p> <p>If you don't enter the credential profile path, the Secure Agent uses the credential profile file present in the following default location in your home directory:</p> <pre>~/.aws/credentials</pre> <p>Note: Database Ingestion and Replication has not been certified with the Credential Profile File Path and Profile Name connection properties. Database Ingestion and Replication finds AWS credentials by using the default credential provider chain that is implemented by the DefaultAWSCredentialsProviderChain class, which includes the credential profile file.</p>
Profile Name	<p>Name of the profile in the credential profile file used to get the credentials.</p> <p>If you don't enter the profile name, the credentials from the default profile in the credential profile file are used.</p>
S3 VPC Endpoint Type	<p>The VPC endpoint type for Amazon S3.</p> <p>You can enable private communication with Amazon S3 by selecting a VPC endpoint. Select one of the following VPC endpoint types:</p> <ul style="list-style-type: none"> - None - Gateway Endpoint - Interface Endpoint <p>Default is None.</p> <p>Doesn't apply to an application ingestion and replication task or database ingestion and replication task.</p>
Endpoint DNS Name for Amazon S3	<p>The DNS name for the Amazon S3 interface endpoint.</p> <p>Enter the DNS name in the following format:</p> <pre>bucket.<DNS name of the interface endpoint></pre> <p>Doesn't apply to an application ingestion and replication task or database ingestion and replication task.</p>
STS VPC Endpoint Type	<p>Applicable when you select the S3 VPC interface endpoint.</p> <p>The VPC endpoint type for AWS STS.</p> <p>When you select IAM Role ARN or Federated SSO IdP, configure the STS VPC endpoint.</p> <p>Doesn't apply to an application ingestion and replication task, database ingestion and replication task, or streaming ingestion and replication task.</p>
Endpoint DNS Name for AWS STS service	<p>The DNS name for the AWS STS interface endpoint.</p> <p>Doesn't apply to an application ingestion and replication task or database ingestion and replication task.</p>
KMS VPC Endpoint Type	<p>Applicable when you select the interface endpoint.</p> <p>The VPC endpoint type for the AWS KMS.</p> <p>Doesn't apply to an application ingestion and replication task or database ingestion and replication task.</p>
Endpoint DNS Name for AWS KMS service	<p>The DNS name for the AWS KMS interface endpoint.</p> <p>Doesn't apply to an application ingestion and replication task or database ingestion and replication task.</p>

Federated user single sign-on connection properties

Configure the following properties when you select ADFS 3.0 in Federated SSO IdP:

Property	Description
Federated User Name	User name of the federated user to access the AWS account through the identity provider.
Federated User Password	Password for the federated user to access the AWS account through the identity provider.
IdP SSO URL	Single sign-on URL of the identity provider for AWS. Doesn't apply to a streaming ingestion and replication task.
SAML Identity Provider ARN	ARN of the SAML identity provider that the AWS administrator created to register the identity provider as a trusted provider.
Role ARN	ARN of the IAM role assumed by the federated user.

Credential Profile File Authentication

You can provide the credentials required to establish the connection with Amazon S3 through the credential profile file that contains an access key and secret key. The credential profile file contains an access key, a secret key, and a session token when you use temporary security credentials.

You can use permanent IAM credentials or temporary security credentials with a session token when you use credential profile file authentication.

If you do not specify the credential profile file path, the default credential file path is used. If you do not specify the profile name, the credentials are used from the default profile in the credential file.

Consider the following rules for a credential profile file:

- The credential file must be on the same machine where you installed the Secure Agent.
- The credential profile file name must end with `.credentials`.
- If you do not specify the credential profile path, the Secure Agent uses the credential profile file present in the following default location in your home directory:

`~/.aws/credentials`

Note: On Windows, you can refer to your home directory by using the environment variable `%UserProfile` `%`. On Unix-like systems, you can use the environment variable `$HOME`.

A sample credential profile file:

```
[default]

aws_access_key_id = 1233333
aws_secret_access_key = abcabcabc

[test-profile]

aws_access_key_id = 1233333
aws_secret_access_key = abcabcabc
aws_session_token = jahaheieomdrftflmlioerp
```

The `aws_access_key_id` and `aws_secret_access_key` specify the AWS access key and secret key used as part of credentials to authenticate the user.

The `aws_session_token` specifies an AWS session token used as part of the credentials to authenticate the user. A session token is required only if you specify temporary security credentials.

Private communication with Amazon S3

You can enable private communication with Amazon S3 by configuring a gateway endpoint or interface endpoint on AWS console and in the Amazon S3 V2 connection.

You can configure Amazon S3 V2 Connector to establish private communication with Amazon S3 without exposing your traffic to the public internet. To access Amazon S3, ensure that the Secure Agent is a part of the subnet in the AWS Virtual Private Cloud (VPC). AWS S3 VPC endpoint enables an S3 request to be routed to the Amazon S3 service, without having to connect a subnet to an internet gateway. You can create an interface endpoint or a gateway endpoint.

For more information, see

[Configuring private communication with Amazon S3 using the Amazon S3 V2 Connector.](#)

AMQP connection properties

When you set up an AMQP connection, you must configure the connection properties.

The following table describes the AMQP connection properties:

Property	Description
Connection Name	Name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
Description	Optional. Description that you can use to identity the connection. The description cannot exceed 4,000 characters.
Type	The AMQP connection type. If you do not see the connection type, go to the Add-On Connectors page to install the connector.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Host Name	Network address of the AMQP broker.
Port	Port number of the AMQP broker to which the underlying TCP connection is made. Default is 5672.
Virtual Host	Virtual host name that identifies the AMQP system. Use the virtual host name for enhanced security.
Username	Username for the AMQP broker.

Property	Description
Password	Password for the AMQP broker.
Use SSL	Enable this option to use SSL for secure transmission. If you enable the SSL authentication, ensure that you provide both keystore and truststore details for using the AMQP connection in a streaming ingestion and replication task.
Keystore File Name	Contains the keys and certificates required for secure communication.
Keystore Password	Password for the keystore filename.
Keystore Type	Type of keystore that you want to use. Keystore type defines the storage and data format of the keystore information and the algorithms used to protect private keys in the keystore. Use one of the following types: <ul style="list-style-type: none"> - JKS. Stores private keys and certificates. - PKCS12. Stores private keys, secret keys, and certificates.
Truststore File Name	Name of the truststore file.
Truststore Password	Password for the truststore file.
Truststore Type	Type of truststore that you want to use. Use one of the following types: <ul style="list-style-type: none"> - JKS - PKCS 12
TLS Protocol	Transport protocols that you want to use. Use one of the following types: <ul style="list-style-type: none"> - SSL - SSLv2Hello - SSLv3 - TLS - TLSv1 - TLSv1.1 - TLSv1.2
Client Authentication	Client authentication policy when connecting to the secured AMQP broker. Use one of the following property values when you define and enable an SSL context. <ul style="list-style-type: none"> - WANT - REQUIRED - NONE

Business 360 Events connection properties

When you create the Business 360 Events connection, you must configure the connection properties.

The following table describes the Business 360 Events connection properties:

Property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 100 characters, contain spaces, or contain the following special characters:~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Type	The connection type. Select Business 360 Events .
Runtime Environment	The name of the runtime environment where you want to run the mappings. Specify a Secure Agent, Hosted Agent, or a serverless runtime environment.
Start Timestamp	A system-generated timestamp variable to set the start of a time range for which you want to get events from the Business 360 data store. Note: You can't modify this attribute.
End Timestamp	A system-generated timestamp variable to set the end of a time range for which you want to get events from the Business 360 data store. Note: You can't modify this attribute.

Cloud Integration Hub connection properties

You can view the Cloud Integration Hub connection only if your organization is provisioned with Cloud Integration Hub. Do not edit, modify, or delete this connection. Do not modify any connection property apart from the **Do not use intermediate staging for subscription flows** and **Use JDBC for Private Publication Repository** properties.

The following table describes the connection properties for a Cloud Integration Hub connection:

Connection Property	Description	Editable
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /	Do not edit.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.	Yes
Type	Cloud Integration Hub connection type.	Do not edit.

Connection Property	Description	Editable
Enable Secret Vault	<p>Stores the publication repository password for the connection in the Secret Manager in the runtime environment that is configured for your organization.</p> <p>This property appears only if secrets manager is set up for your organization.</p> <p>Select this option to use the credentials from the Secret Manager. If you don't enable this option, the credentials are stored in the repository or on a local Secure Agent, depending on how your organization is configured.</p> <p>For information about how to configure and use a secrets manager, see "Secrets manager configuration" in the Administrator help.</p>	Do not edit.
Runtime Environment	Name of the runtime environment where you want to run the tasks.	Do not edit.
Do not use intermediate staging for subscription flows	Disables writing to intermediate staging. Enable this property if you do not want to write to the intermediate staging, the Data Integration task reads the data from Cloud Integration Hub and then writes the data directly to the target location. Disabling writing to intermediate staging might affect system performance.	Yes
Use JDBC for Private Publication Repository	<p>To configure zero downtime for a private publication repository. Enable this property to ensure uninterrupted access to data on the private publication repository. You can enable zero downtime for publications and subscriptions that trigger a Data Integration task.</p> <p>On a hosted publication repository, Cloud Integration Hub applies zero downtime by default for all publication and subscription types.</p>	Yes

Databricks connection properties

Create a Databricks connection to securely read data from or write data to Databricks.

Staging prerequisites

Before you create a connection, you must perform certain prerequisite tasks to configure the staging environment to connect to SQL warehouse or Databricks cluster.

SQL warehouse

Configure either the AWS or Azure staging environment for the SQL warehouse based on the deployed environment. You also need to configure the Spark parameters for the SQL warehouse to use Azure and AWS staging.

You can use a SQL warehouse on the Windows and Linux operating systems.

Configure AWS staging

Configure IAM AssumeRole authentication to use AWS staging for the SQL warehouse.

Configure Spark parameters for AWS staging

On the Databricks SQL Admin console, navigate to **SQL Warehouse Settings > Data Security**, and then configure the Spark parameters for AWS under **Data access configuration**.

Add the following Spark configuration parameters and restart the SQL warehouse:

- `spark.hadoop.fs.s3a.access.key` <S3 Access Key value>
- `spark.hadoop.fs.s3a.secret.key` <S3 Secret Key value>
- `spark.hadoop.fs.s3a.endpoint` <S3 Staging Bucket endpoint value>

For example, the S3 staging bucket warehouse value is `s3.ap-south-1.amazonaws.com`.

Ensure that the configured access key and secret key have access to the S3 buckets where you store the data for Databricks tables.

Configure Azure staging

Before you use Microsoft Azure Data Lake Storage Gen2 to stage files, perform the following tasks:

- Create a storage account to use with Microsoft Azure Data Lake Storage Gen2 and enable **Hierarchical namespace** in the Azure portal.
You can use role-based access control to authorize the users to access the resources in the storage account. Assign the Contributor role or Reader role to the users. The contributor role grants you full access to manage all resources in the storage account, but does not allow you to assign roles. The reader role allows you to view all resources in the storage account, but does not allow you to make any changes.
Note: To add or remove role assignments, you must have write and delete permissions, such as an Owner role.
- Register an application in Azure Active Directory to authenticate users to access the Microsoft Azure Data Lake Storage Gen2 account.
You can use role-based access control to authorize the application. Assign the Storage Blob Data Contributor or Storage Blob Data Reader role to the application. The Storage Blob Data Contributor role lets you read, write, and delete Azure Storage containers and blobs in the storage account. The Storage Blob Data Reader role lets you only read and list Azure Storage containers and blobs in the storage account.
- Create an Azure Active Directory web application for service-to-service authentication with Microsoft Azure Data Lake Storage Gen2.
Note: Ensure that you have superuser privileges to access the folders or files created in the application using the connector.
- To read complex files, set the JVM options for type DTM to increase the -Xms and -Xmx values in the system configuration details of the Secure Agent to avoid java heap space error. The recommended -Xms and -Xmx values are 512 MB and 1024 MB respectively.

Configure Spark parameters for Azure staging

On the Databricks SQL Admin console, navigate to **SQL Warehouse Settings > Data Security**, and then configure the Spark parameters for Azure under **Data access configuration**.

Add the following Spark configuration parameters and restart the SQL warehouse:

- `spark.hadoop.fs.azure.account.oauth2.client.id.<storage-account-name>.dfs.core.windows.net <value>`
- `spark.hadoop.fs.azure.account.auth.type.<storage-account-name>.dfs.core.windows.net OAuth`
- `spark.hadoop.fs.azure.account.oauth2.client.secret.<storage-account-name>.dfs.core.windows.net <Value>`
- `spark.hadoop.fs.azure.account.oauth.provider.type.<storage-account-name>.dfs.core.windows.net org.apache.hadoop.fs.azurebfs.oauth2.ClientCredsTokenProvider`
- `spark.hadoop.fs.azure.account.oauth2.client.endpoint.<storage-account-name>.dfs.core.windows.net https://login.microsoftonline.com/<Tenant ID>/oauth2/token`

Ensure that the configured client ID and client secret have access to the file systems where you store the data for Databricks tables.

Databricks cluster

Configure the Spark parameters for the Databricks cluster to use Azure and AWS staging based on where the cluster is deployed.

You also need to enable the Secure Agent properties for runtime and design-time processing on the Databricks cluster.

You can use a Databricks cluster only on the Linux operating system.

Spark configuration

Before you connect to the Databricks cluster, you must configure the Spark parameters on AWS and Azure.

Configuration on AWS

Add the following Spark configuration parameters for the Databricks cluster and restart the cluster:

- `spark.hadoop.fs.s3a.access.key <value>`
- `spark.hadoop.fs.s3a.secret.key <value>`
- `spark.hadoop.fs.s3a.endpoint <value>`

Ensure that the access and secret key configured has access to the buckets where you store the data for Databricks tables.

Configuration on Azure

Add the following Spark configuration parameters for the Databricks cluster and restart the cluster:

- `fs.azure.account.oauth2.client.id.<storage-account-name>.dfs.core.windows.net <value>`
- `fs.azure.account.auth.type.<storage-account-name>.dfs.core.windows.net <value>`
- `fs.azure.account.oauth2.client.secret.<storage-account-name>.dfs.core.windows.net <Value>`
- `fs.azure.account.oauth.provider.type.<storage-account-name>.dfs.core.windows.net org.apache.hadoop.fs.azurebfs.oauth2.ClientCredsTokenProvider`
- `fs.azure.account.oauth2.client.endpoint.<storage-account-name>.dfs.core.windows.net https://login.microsoftonline.com/<Tenant ID>/oauth2/token`

Ensure that the client ID and client secret configured has access to the file systems where you store the data for Databricks tables.

Connect to Databricks

Let's configure the Databricks connection properties to connect to Databricks.

Before you begin

You can use a Databricks connection to connect to SQL warehouse or Databricks cluster to read from and write to Databricks tables.

- **SQL warehouse**

The Secure Agent connects to the SQL warehouse at design time and runtime.

- **Databricks cluster**

The term Databricks cluster refers to the all-purpose cluster and job cluster. The Secure Agent connects to the all-purpose cluster to import the metadata at design time and to the job cluster to run the mappings.

Before you get started, you'll need to get information from your Databricks account.

The following video shows you how to get the information you need:



You also need to configure the AWS or Azure staging environment to use the SQL endpoint or the Databricks cluster in the connection.

To learn about the staging prerequisites for the Azure or AWS environment, check out [“SQL warehouse” on page 57](#) or [“Databricks cluster” on page 59](#).

Connection details

The following table describes the basic connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	Databricks
Runtime Environment	The name of the runtime environment where you want to run tasks. You cannot run an application ingestion and replication, database ingestion and replication, or streaming ingestion and replication task on a Hosted Agent or serverless runtime environment.

Property	Description
SQL Warehouse JDBC URL	<p>Databricks SQL Warehouse JDBC connection URL.</p> <p>Required to connect to a Databricks SQL warehouse. Doesn't apply to Databricks clusters.</p> <p>Note: Databricks SQL Serverless is the recommended Databricks cluster type.</p> <p>To get the SQL Warehouse JDBC URL, go to the Databricks console and select the JDBC driver version from the JDBC URL menu.</p> <p>Application ingestion and replication and database ingestion and replication tasks can use JDBC URL version 2.6.25 or later or 2.6.22 or earlier. The URLs must begin with the prefix <code>jdbc:databricks://</code>, as follows:</p> <pre>jdbc:databricks://<Databricks Host>:443/default;transportMode=http;ssl=1;AuthMech=3;httpPath=/sql/1.0/endpoints/<SQL endpoint cluster ID>;</pre> <p>Ensure that you set the required environment variables in the Secure Agent. Also specify the correct JDBC Driver Class Name under advanced connection settings.</p> <p>Note: Specify the database name in the Database Name connection property. If you specify the database name in the JDBC URL, it is not considered. The Databricks Host, Organization ID, and Cluster ID properties are not considered if you configure the SQL warehouse JDBC URL property.</p>
Databricks Token	<p>Personal access token to access Databricks.</p> <p>Required for SQL warehouse and Databricks cluster.</p>
Catalog Name	<p>If you use Unity Catalog, the name of an existing catalog in the metastore.</p> <p>Optional for SQL warehouse. Doesn't apply to Databricks cluster.</p> <p>You can also specify the catalog name in the end of the SQL warehouse JDBC URL.</p> <p>Note: The catalog name cannot contain special characters.</p> <p>For more information about Unity Catalog, see the Databricks documentation.</p>

Advanced settings

The following table describes the advanced connection properties:

Property	Description
Database	<p>The name of the schema in the unity catalog in Databricks.</p> <p>Optional for SQL warehouse and Databricks cluster.</p> <p>If you do not specify a value, all databases available in the workspace are listed. The value you specify overrides the schema specified in the SQL Warehouse JDBC URL connection property.</p>
JDBC Driver Class Name	<p>The name of the JDBC driver class.</p> <p>Optional for SQL warehouse and Databricks cluster.</p> <p>For JDBC URL versions 2.6.22 or earlier, specify the driver class name as <code>com.simba.spark.jdbc.Driver</code>.</p> <p>For JDBC URL versions 2.6.25 or later, specify the driver class name as <code>com.databricks.client.jdbc.Driver</code>.</p>

Property	Description
Staging Environment	<p>The cloud provider where the Databricks cluster is deployed. Required for SQL warehouse and Databricks cluster.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - AWS - Azure - Personal Staging Location <p>Default is Personal Staging Location.</p> <p>You can select the Personal Staging Location as the staging environment instead of Azure or AWS staging environments to stage data locally for mappings and tasks.</p> <p>If you select Personal Staging Location for a connection that Data Ingestion and Replication uses, the Parquet data files for application ingestion and replication jobs or database ingestion and replication jobs can be staged to a local personal storage location, which has a data retention period of 7 days. You must also specify a Database Host value. If you use Unity Catalog, note that a personal storage location is automatically provisioned.</p> <p>Personal staging location doesn't apply to Databricks cluster.</p> <p>You cannot use personal staging location with Databricks unmanaged tables.</p> <p>Note: You cannot switch between clusters once you establish a connection.</p>
Databricks Host	<p>The host name of the endpoint the Databricks account belongs to. Required for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>You can get the Databricks Host from the JDBC URL. The URL is available in the Advanced Options of JDBC or ODBC in the Databricks all-purpose cluster.</p> <p>The following example shows the Databricks Host in JDBC URL:</p> <pre>jdbc:spark://<Databricks Host>:443/ default;transportMode=http; ssl=1;httpPath=sql/protocolv1/o/<Org Id>/<Cluster ID>; AuthMech=3; UID=token; PWD=<personal-access-token></pre> <p>The value of PWD in Databricks Host, Organization Id, and Cluster ID is always <personal-access-token>.</p>
Cluster ID	<p>The ID of the cluster. Required for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>You can get the cluster ID from the JDBC URL. The URL is available in the Advanced Options of JDBC or ODBC in the Databricks all-purpose cluster</p> <p>The following example shows the Cluster ID in JDBC URL:</p> <pre>jdbc:spark://<Databricks Host>:443/ default;transportMode=http; ssl=1;httpPath=sql/protocolv1/o/<Org Id>/<Cluster ID>; AuthMech=3; UID=token; PWD=<personal-access-token></pre>
Organization ID	<p>The unique organization ID for the workspace in Databricks. Required for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>You can get the Organization ID from the JDBC URL. The URL is available in the Advanced Options of JDBC or ODBC in the Databricks all-purpose cluster</p> <p>The following example shows the Organization ID in JDBC URL:</p> <pre>jdbc:spark://<Databricks Host>:443/ default;transportMode=http; ssl=1;httpPath=sql/protocolv1/o/<Organization ID>/ <Cluster ID>;AuthMech=3;UID=token; PWD=<personal-access-token></pre>
Min Workers	<p>The minimum number of worker nodes to be used for the Spark job. Minimum value is 1. Required for Databricks cluster. Doesn't apply to SQL warehouse.</p>
Max Workers	<p>The maximum number of worker nodes to be used for the Spark job. If you don't want to autoscale, set Max Workers = Min Workers or don't set Max Workers. Optional for Databricks cluster. Doesn't apply to SQL warehouse.</p>

Property	Description
DB Runtime Version	<p>The version of Databricks cluster to spawn when you connect to Databricks cluster to process mappings.</p> <p>Required for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>Select the Databricks runtime version 9.1 LTS or 13.3 LTS.</p>
Worker Node Type	<p>The worker node instance type that is used to run the Spark job.</p> <p>Required for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>For example, the worker node type for AWS can be i3.2xlarge. The worker node type for Azure can be Standard_DS3_v2.</p>
Driver Node Type	<p>The driver node instance type that is used to collect data from the Spark workers.</p> <p>Optional for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>For example, the driver node type for AWS can be i3.2xlarge. The driver node type for Azure can be Standard_DS3_v2.</p> <p>If you don't specify the driver node type, Databricks uses the value you specify in the worker node type field.</p>
Instance Pool ID	<p>The instance pool ID used for the Spark cluster.</p> <p>Optional for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>If you specify the Instance Pool ID to run mappings, the following connection properties are ignored:</p> <ul style="list-style-type: none"> - Driver Node Type - EBS Volume Count - EBS Volume Type - EBS Volume Size - Enable Elastic Disk - Worker Node Type - Zone ID
Elastic Disk	<p>Enables the cluster to get additional disk space.</p> <p>Optional for Databricks cluster. Doesn't apply to SQL warehouse.</p> <p>Enable this option if the Spark workers are running low on disk space.</p>
Spark Configuration	Doesn't apply to a data loader task or to Data Ingestion and Replication tasks.
Spark Environment Variables	Doesn't apply to a data loader task or to Data Ingestion and Replication tasks.

AWS staging environment

The following table describes the properties for the AWS staging environment:

Property	Description
S3 Authentication Mode	The authentication mode to access Amazon S3. Select one of the following authentication modes: <ul style="list-style-type: none">- Permanent IAM credentials. Uses the S3 access key and S3 secret key to connect to Databricks.- IAM Assume Role. Uses the AssumeRole for IAM authentication to connect to Databricks. Doesn't apply to Databricks cluster.
S3 Access Key	The key to access the Amazon S3 bucket.
S3 Secret Key	The secret key to access the Amazon S3 bucket.
S3 Data Bucket	The existing bucket to store the Databricks data.
S3 Staging Bucket	The existing bucket to store the staging files.
S3 VPC Endpoint Type	The type of Amazon Virtual Private Cloud endpoint for Amazon S3. You can use a VPC endpoint to enable private communication with Amazon S3. Select one of the following options: <ul style="list-style-type: none">- None. Select if you do not want to use a VPC endpoint.- Interface Endpoint. Select to establish private communication with Amazon S3 through an interface endpoint which uses a private IP address from the IP address range of your subnet. It serves as an entry point for traffic destined to an AWS service.
Endpoint DNS Name for S3	The DNS name for the Amazon S3 interface endpoint. Replace the asterisk symbol with the bucket keyword in the DNS name. Enter the DNS name in the following format: <code>bucket.<DNS name of the interface endpoint></code> For example, <code>bucket.vpce-s3.us-west-2.vpce.amazonaws.com</code>
IAM Role ARN	The Amazon Resource Number (ARN) of the IAM role assumed by the user to use the dynamically generated temporary security credentials. Set the value of this property if you want to use the temporary security credentials to access the Amazon S3 staging bucket. For more information about how to get the ARN of the IAM role, see the <i>AWS documentation</i> .
Use EC2 Role to Assume Role	Optional. Select the check box to enable the EC2 role to assume another IAM role specified in the IAM Role ARN option. The EC2 role must have a policy attached with a permission to assume an IAM role from the same or different AWS account.
STS VPC Endpoint Type	The type of Amazon Virtual Private Cloud endpoint for AWS Security Token Service. You can use a VPC endpoint to enable private communication with Amazon Security Token Service. Select one of the following options: <ul style="list-style-type: none">- None. Select if you do not want to use a VPC endpoint.- Interface Endpoint. Select to establish private communication with Amazon Security Token Service through an interface endpoint which uses a private IP address from the IP address range of your subnet.

Property	Description
Endpoint DNS Name for AWS STS	The DNS name for the AWS STS interface endpoint. For example, <code>vpce-01f22cc14558c241f-s8039x4c.sts.us-west-2.vpce.amazonaws.com</code>
S3 Service Regional Endpoint	The S3 regional endpoint when the S3 data bucket and the S3 staging bucket need to be accessed through a region-specific S3 regional endpoint. Doesn't apply to Databricks cluster. Default is <code>s3.amazonaws.com</code> .
S3 Region Name	The AWS cluster region in which the bucket you want to access resides. Select a cluster region if you choose to provide a custom JDBC URL that does not contain a cluster region name in the JDBC URL connection property.
Zone ID	The zone ID for the Databricks job cluster. Optional for Databricks cluster. Doesn't apply to SQL warehouse. Applies only if you want to create a Databricks job cluster in a particular zone at runtime. For example, <code>us-west-2a</code> . Note: The zone must be in the same region where your Databricks account resides.
EBS Volume Type	The type of EBS volumes launched with the cluster. Optional for Databricks cluster. Doesn't apply to SQL warehouse.
EBS Volume Count	The number of EBS volumes launched for each instance. You can choose up to 10 volumes. Optional for Databricks cluster. Doesn't apply to SQL warehouse. Note: In a Databricks connection, specify at least one EBS volume for node types with no instance store. Otherwise, cluster creation fails.
EBS Volume Size	The size of a single EBS volume in GiB launched for an instance. Optional for Databricks cluster. Doesn't apply to SQL warehouse.

Azure staging environment

The following table describes the properties for the Azure staging environment:

Property	Description
ADLS Storage Account Name	The name of the Microsoft Azure Data Lake Storage account.
ADLS Client ID	The ID of your application to complete the OAuth Authentication in the Active Directory.
ADLS Client Secret	The client secret key to complete the OAuth Authentication in the Active Directory.
ADLS Tenant ID	The ID of the Microsoft Azure Data Lake Storage directory that you use to write data.
ADLS Endpoint	The OAuth 2.0 token endpoint from where authentication based on the client ID and client secret is completed.

Property	Description
ADLS Filesystem Name	The name of an existing file system to store the Databricks data.
ADLS Staging Filesystem Name	The name of an existing file system to store the staging data.

JDBC URL parameters

You can utilize the additional JDBC URL parameters field in the Databricks connection to customize and set any additional parameters required to connect to Databricks.

You can configure the following properties as additional JDBC URL parameters in the Databricks connection:

- To pass the unity catalog information to Databricks, specify the catalog name after the SQL warehouse cluster ID in the following format:

```
jdbc:spark://<Databricks Host>:443/
default;transportMode=http;ssl=1;AuthMech=3;httpPath=/sql/1.0/endpoints/<SQL endpoint
cluster ID>;ConnCatalog=<catalog_name>;
```

- To connect to Databricks using the proxy server, enter the following parameters:

```
jdbc: spark://<Databricks Host>:443/
default;transportMode=http;ssl=1;AuthMech=3;httpPath=/sql/1.0/warehouses/
219fe3013963cdce;UseProxy=<Proxy=true>;ProxyHost=<proxy host IPaddress>;ProxyPort=<proxy
server port number>;ProxyAuth=<Auth_true>;
```

Note: Data Ingestion and Replication does not support use of a proxy server to connect to Databricks.

- To connect to SSL-enabled Databricks, specify the value in the JDBC URL in the following format:

```
jdbc:spark://<Databricks Host>:443/
default;transportMode=http;ssl=1;AuthMech=3;httpPath=/sql/1.0/endpoints/<SQL endpoint
cluster ID>;
```

Rules and guidelines for personal staging location

When you select the personal staging location as a staging environment, the data is first staged in a java temporary location and then copied to a personal staging location of the unity catalog. Both the staged files will be deleted after the task runs successfully.

However, to stage the data in a different directory, configure the DTM property `-Djava.io.tmpdir=/my/dir/path` in the JVM options in the system configuration settings of the Administrator service.

To enable data staging in a different directory, you should have read and write permission and enough disk space to stage the data in the directory.

When you specify a personal staging location in the Databricks connection properties for staging, consider the following rules and guidelines:

- You can only specify unity enabled catalog in the SQL warehouse JDBC URL.
- All mappings that are configured run without SQL ELT optimization.
- The data is staged in the folder `stage://tmp/<user_name>` where the `<user_name>` is picked from the Databricks token provided in the connection and requires read and write access to the personal staging location in root location of AWS and Azure.

Db2 for i Database Ingestion connection properties

When you define a Db2 for i Database Ingestion connection, you must configure connection properties. You can use this connection type in database ingestion and replication tasks, which you configure in the Data Ingestion and Replication service.

The following table describes the connection properties:

Property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Ensure that the type is Db2 for i Database Ingestion .
Runtime Environment	The name of the runtime environment where you want to run database ingestion and replication tasks. You define runtime environments in Administrator.
User Name	The user name to use for connecting to the Db2 for i instance.
Password	The password to use for connecting to the Db2 for i instance.
Host	The name of the machine that hosts the database server.
Port	The network port number used to connect to the database server.
Location Name	The name of the Db2 for i location that you want to access. Your system administrator can determine the name of the Db2 location by using the WRKRDBDIRE command. In the output, find the name of the database that is listed as *LOCAL and then use that value as the value of this property.
JDBC Driver	The type of JDBC driver. Select one of the following options: - Data Direct - JTOpen Default is Data Direct.
Code Page for Bit Data	The code page that Database Ingestion and Replication uses to read character data that is stored as bit data. This value must be a canonical name for the java.io API and java.lang API. For more information, see the supported encodings in the Oracle Java documentation. Specify this property if you have FOR BIT DATA source columns.

Property	Description
Advanced Connection Properties	<p>Advanced properties for the JDBC driver which is used to connect to the Db2 for i source. If you specify more than one <i>property=value</i> entry, separate them with a semicolon (;).</p> <p>For information about the DataDirect JDBC driver connection properties, see Progress DataDirect documentation. For example, you can set the ConnectionRetryCount property to control the number of times the driver retries attempts to connect to the primary database server.</p> <p>For information about the JTOpen JDBC driver connection properties, see IBM Toolbox for Java JDBC properties.</p>
Encryption Method	<p>The data encryption method for the JTOpen JDBC Driver.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - No Encryption - SSL <p>Default is No Encryption.</p> <p>If you select SSL, you must add the required certificates to the Informatica Cloud Secure Agent JRE cacerts keystore in one of the following locations:</p> <p>For Linux:</p> <pre>Secure Agent Directory\jdk\jre\lib\security\cacerts</pre> <p>For Windows:</p> <pre>Secure Agent Directory\apps\jdkLatestVersion\jre</pre>

Db2 for LUW Database Ingestion connection properties

When you define a Db2 for LUW Database Ingestion connection, you must configure connection properties. You can use this connection type in database ingestion and replication tasks, which you configure in the Data Ingestion and Replication service.

The following table describes the connection properties:

Property	Description
Connection Name	<p>A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -</p> <p>Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.</p>
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Ensure that the type is Db2 for LUW Database Ingestion .
Runtime Environment	The name of the runtime environment where you want to run database ingestion and replication tasks. You define runtime environments in Administrator.
User Name	The user name to use for connecting to the Db2 for LUW instance.
Password	The password to use for connecting to the Db2 for LUW instance.

Property	Description
Host	The name of the machine that hosts the database server.
Port	The network port number used to connect to the database server.
Database Name	The name of the Db2 for LUW database that you want to access.
Advanced Connection Properties	<p>Advanced properties for the Progress DataDirect JDBC DB2 driver, which is used to connect to the Db2 for LUW source. If you specify more than one <i>property=value</i> entry, separate them with a semicolon (;).</p> <p>The driver properties that you can enter in this field are described in the Progress DataDirect for JDBC connection properties. For example, you can set the EncryptionMethod property to control whether data is encrypted and decrypted when transmitted over the network between the driver and database server.</p>

Db2 for zOS Database Ingestion connection properties

When you define a Db2 for zOS Database Ingestion connection, you must configure connection properties. You can use this connection type in database ingestion and replication tasks, which you configure in the Data Ingestion and Replication service.

The following table describes the connection properties:

Property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Ensure that the type is Db2 for zOS Database Ingestion .
Runtime Environment	The name of the runtime environment where you want to run database ingestion and replication tasks. You define runtime environments in Administrator.
User Name	The user name to use for connecting to the Db2 for z/OS instance.
Password	The password to use for connecting to the Db2 for z/OS instance.
Host	The name of the machine that hosts the database server.
Port	The network port number used to connect to the database server.
Location Name	The name of the Db2 for z/OS location that you want to access. For Db2 for z/OS, your system administrator can determine the name of your Db2 location using the command DISPLAY DDF.

Property	Description
Code Page for Bit Data	The code page that Database Ingestion and Replication uses to read character data that is stored as bit data. This value must be a canonical name for the java.io API and java.lang API. For more information, see the supported encodings in the Oracle Java documentation. Specify this property if you have FOR BIT DATA source columns.
CDC Stored Procedure Schema	For incremental change data capture processing, the name of the schema for the z/OS stored procedure that is required to collect change data from the Db2 log. This value is specified in the #STPINST data set that you customized when setting up the stored procedure on z/OS. No default value is provided.
CDC Stored Procedure Name	For incremental change data capture processing, the name of the z/OS stored procedure that is required to collect change data from the Db2 log. This value is specified in the #STPINST data set that you customized when setting up the stored procedure on z/OS. The default value is INFALOG.
Advanced Connection Properties	<p>Advanced properties for the Progress DataDirect JDBC Db2 driver, which is used to connect to the Db2 for z/OS source. If you specify more than one <i>property=value</i> entry, separate them with a semicolon (;).</p> <p>The driver properties that you can enter in this field are described in the Progress DataDirect documentation at https://docs.progress.com/bundle/datadirect-connect-jdbc-51/page/Connection-Properties_10.html. For example, you can set the ConnectionRetryCount property to control the number of times the driver retries attempts to connect to the primary database server.</p>

Flat file connection properties

Defines the properties you need to assign to for a flat file source connection.

The following table describes the flat file connection properties:

Connection Property	Description
Connection Name	<p>A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -</p> <p>Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.</p>
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Select the Flat File connection type.
Runtime Environment	<p>Runtime environment that contains the Secure Agent to use for accessing the flat files.</p> <p>Note: Do not select a runtime environment with Secure Agents that run on NTT. A flat file connection cannot use a Secure Agent that runs on NTT.</p>

Connection Property	Description
Directory	<p>Directory where the flat file is stored. Must be accessible by all Secure Agents in the selected runtime environment.</p> <p>Enter the full directory or click Browse to locate and select the directory.</p> <p>When you use the connection, you can select a file that's contained in the directory or in any of its subdirectories.</p> <p>Maximum length is 100 characters. Directory names can contain alphanumeric characters, spaces, and the following special characters:</p> <p>/ \ : _ ~</p> <p>The directory is the service URL for this connection type.</p> <p>Note: On Windows, the Browse for Directory dialog box doesn't display mapped drives. You can browse My Network Places in Windows Explorer to locate the directory and copy the location from the address bar or enter the directory name in the following format: \<server_name>\<directory_path>. If network directories do not display, you can configure a login for the Secure Agent service. This functionality might not be available on newer versions of Windows.</p> <p>Do not include the name of the flat file. You specify the file name when you create the task.</p>
Browse button	Use to locate and select the directory where flat files are stored.
Date Format	<p>Date format for date fields in the flat file. Default date format is:</p> <p>MM/dd/yyyy HH:mm:ss</p>

Connection Property	Description
Code Page	<p>The code page of the system that hosts the flat file. Select one of the following code pages:</p> <ul style="list-style-type: none"> - MS Windows Latin 1. Select for ISO 8859-1 Western European data - UTF-8. Select for Unicode data - UTF-16 encoding of Unicode (Big Endian) - UTF-16 encoding of Unicode (Lower Endian) - Shift-JIS. Select for double-byte character data. - ISO 8859-15 Latin 9 (Western European) - ISO 8859-2 Eastern European - ISO 8859-3 Southeast European - ISO 8859-5 Cyrillic - ISO 8859-9 Latin 5 (Turkish) - IBM EBCDIC International Latin-1 - Japanese EUC (with \ <-> Yen mapping) - IBM EBCDIC Japanese - IBM EBCDIC Japanese CP939 - PC Japanese SJIS-78 syntax (IBM-942) - PC Japanese SJIS-90 (IBM-943) - MS Windows Traditional Chinese, superset of Big 5 - Taiwan Big-5 (w/o euro update) - Chinese EUC - ISO 8859-8 Hebrew - PC Hebrew (old) - PC Hebrew (w/o euro update) - EBCDIC Hebrew (updated with new sheqel, control characters) - IBM EBCDIC US English IBM037 - UTF-32 encoding of Unicode (Lower Endian) - ISO 8859-1 Western European. - IBM EBCDIC French - ISO 8859-10 Latin 6 (Nordic) * - EBCDIC Finland, Sweden - MOS-DOS Thai, superset of TIS 620 - 7-bit ASCII - EBCDIC Finland, Sweden (w/euro update) - MS-DOS Windows Latin 2 (Central Europe) - Japanese EBCDIC-Kana Fujitsu <p>In advanced mappings, flat file objects in cloud storage connections must use UTF-8 encoding.</p> <p>If the file contains supplementary characters with UTF-16 encoding, the task fails.</p> <p>Note: When you use a flat file connection with the Shift-JIS code page and a UTF data object, be sure to install fonts that fully support Unicode.</p>
<p>* Data preview uses a similar ISO 8859-4 Scandinavian/Baltic code page, but runtime processing uses ISO 8859-10 Latin 6 (Nordic), so data preview and runtime encoding won't match.</p>	

Google Analytics Mass Ingestion connection properties

When you set up a Google Analytics Mass Ingestion connection, you must configure the connection properties.

The following table describes the connection properties for a Google Analytics Mass Ingestion connection:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Select the Google Analytics Mass Ingestion connection type.
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Service Account ID	The client_email value in the JSON file that you download after you create a service account.
Service Account Key	The private_key value in the JSON file that you download after you create a service account.

Google BigQuery V2 connection properties

When you create a Google BigQuery V2 connection, configure the connection properties.

The following table describes the Google BigQuery V2 connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Google BigQuery V2 connection type.

Property	Description
Runtime Environment	Name of the runtime environment where you want to run the tasks. You cannot run a database ingestion and replication task on a Hosted Agent or serverless runtime environment.
Service Account Email	The client_email value in the JSON file that you download after you create a service account.
Service Account Key	The private_key value in the JSON file that you download after you create a service account.
Project ID	The project_id value in the JSON file that you download after you create a service account. If you have created multiple projects with the same service account, enter the ID of the project that contains the dataset that you want to connect to.
Storage Path	Path in Google Cloud Storage where the agent creates a local stage file to store the data temporarily. Applies to tasks that read or write large volumes of data. Use this property when you read data in staging mode or write data in bulk mode. You can either enter the bucket name or the bucket name and folder name. Use one of the following formats: <ul style="list-style-type: none"> - gs://<bucket_name> - gs://<bucket_name>/<folder_name>
Connection mode	The mode that you want to use to read data from or write data to Google BigQuery. Select one of the following connection modes: <ul style="list-style-type: none"> - Simple. Flattens each field within the Record data type field as a separate field in the mapping. - Hybrid. Displays all the top-level fields in the Google BigQuery table including Record data type fields. Google BigQuery V2 Connector displays the top-level Record data type field as a single field of the String data type in the mapping. - Complex. Displays all the columns in the Google BigQuery table as a single field of the String data type in the mapping. Default is Simple.
Schema Definition File Path	Directory on the Secure Agent machine where the Secure Agent must create a JSON file with the sample schema of the Google BigQuery table. The JSON file name is the same as the Google BigQuery table name. Alternatively, you can specify a storage path in Google Cloud Storage where the Secure Agent must create a JSON file with the sample schema of the Google BigQuery table. You can download the JSON file from the specified storage path in Google Cloud Storage to a local machine. The schema definition file is required if you configure complex connection mode in the following scenarios: <ul style="list-style-type: none"> - You add a Hierarchy Builder transformation in a mapping to read data from relational sources and write data to a Google BigQuery target. - You add a Hierarchy Parser transformation in a mapping to read data from a Google BigQuery source and write data to relational targets.
Use Legacy SQL For Custom Query	Select this option to use a legacy SQL to define a custom query. If you clear this option, you must use a standard SQL to define a custom query. Note: Not applicable when you configure the Google BigQuery V2 connection in hybrid or complex mode.
Dataset Name for Custom Query	When you define a custom query, you must specify a Google BigQuery dataset.

Property	Description
Region ID	The region name where the Google BigQuery dataset that you want to access resides. Note: You must ensure that you specify a bucket name or the bucket name and folder name in the Storage Path property that resides in the specified region. For more information about the regions supported by Google BigQuery, see Dataset locations .
Staging Dataset	The Google BigQuery dataset name where you want to create the staging table to stage the data. You can define a Google BigQuery dataset that is different from the source or target dataset.
Provide Optional Properties	Comma-separated key-value pairs of custom properties in the Google BigQuery V2 connection to configure certain source and target functionalities. For more information about the list of custom properties that you can specify, see the Informatica Knowledge Base article: https://kb.informatica.com/faq/7/Pages/26/632722.aspx

Note: Ensure that you specify valid credentials in the connection properties. The test connection is successful even if you specify incorrect credentials in the connection properties.

Google Cloud Storage V2 connection properties

When you create a Google Cloud Storage V2 connection, configure the connection properties.

The following table describes the Google Cloud Storage connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Google Cloud Storage V2 connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks. You cannot run a database ingestion and replication task or streaming ingestion and replication task on a Hosted Agent or serverless runtime environment.
Service Account Email	The client_email value in the JSON file that you download after you create a service account.
Service Account Key	The private_key value in the JSON file that you download after you create a service account.
Project ID	The project_id value in the JSON file that you download after you create a service account. If you created multiple projects with the same service account, enter the ID of the project that contains the bucket that you want to connect to.

Property	Description
Private Key ID	The <code>private_key_id</code> value in the JSON file that you download after you create a service account. This property applies only to a database ingestion and replication or streaming ingestion and replication task.
Client ID	The <code>client_id</code> value in the JSON file that you download after you create a service account. This property applies only to a database ingestion and replication or streaming ingestion and replication task.
Bucket Name	The Google Cloud Storage bucket name that you want to connect to. When you select a source object, the Package Explorer lists files and folder available in the specified Google Cloud Storage bucket. If you do not specify a bucket name, you can select a bucket from the Package Explorer to select a source.

Configuring the proxy settings on Windows

If your organization uses an outgoing proxy server to connect to the Internet, the Secure Agent connects to Informatica Intelligent Cloud Services using the proxy server.

To configure the proxy server settings for the Secure Agent on a Windows machine, you must configure the proxy server settings through the Secure Agent Manager and the JVM options of the Secure Agent.

Restriction: These steps do not work for Data Ingestion and Replication.

Contact your network administrator for the proxy settings.

1. Click **Start > All Programs > Informatica Cloud Secure Agent > Informatica Cloud Secure Agent** to launch the Secure Agent Manager.
The **Secure Agent Manager** displays the **Secure Agent** status.
2. Click **Proxy** on the Secure Agent Manager page.
3. Click **Use a Proxy Server** to enter the proxy server settings.
4. Configure the following proxy server details:

Field	Description
Proxy Host	Host name of the outgoing proxy server that the Secure Agent uses.
Proxy Port	Port number of the outgoing proxy server.
User Name	User name to connect to the outgoing proxy server.
Password	Password to connect to the outgoing proxy server.

5. Click **OK**.
6. Log in to Informatica Intelligent Cloud Services.
7. Open Administrator and select **Runtime Environments**.
8. Select the Secure Agent for which you want to configure a proxy server.
9. On the upper-right corner of the page, click **Edit**.

10. In the **System Configuration Details** section, select the **Type** as **Agent** for the CMI Streaming Agent Service.
11. To use a proxy server, add the following parameters in any **JVMOption** field and specify appropriate values for each parameter:

Parameter	Description
-Dproxy.host=	Host name of the outgoing HTTPS proxy server.
-Dproxy.port=	Port number of the outgoing HTTPS proxy server.
-Dproxy.user=	User name for the HTTPS proxy server.
-Dproxy.password=	Password for the HTTPS proxy server.

Note: You must specify the parameter and the value for the parameter enclosed in single quotation marks.

For example,

```
JVMOption1='-Dproxy.host=INPQ8583WI29'
```

```
JVMOption2='-Dproxy.port=8081'
```

```
JVMOption3='-Dproxy.user=adminuser'
```

```
JVMOption4='-Dproxy.password=password'
```

Note: You can configure only five **JVMOption** fields in the **System Configuration Details** section. To configure the remaining parameters, you must add the **JVMOption** fields in the **Custom Configuration Details** section. In the **Custom Configuration Details** section, select the **Type** as **Agent** for the CMI Streaming Agent Service, add the **JVMOption** fields, and specify the remaining parameters and appropriate values for each parameter.

12. Click **Save**.

The Secure Agent restarts to apply the settings.

Note: The session log does not record the proxy server details even if you have configured a proxy server.

Configuring the proxy settings on Linux

If your organization uses an outgoing proxy server to connect to the Internet, the Secure Agent connects to Informatica Intelligent Cloud Services through the proxy server.

You can update the proxy server settings defined for the Secure Agent from the command line. To configure the proxy server settings for the Secure Agent on a Linux machine, you must update the `proxy.ini` file and configure the JVM options of the Secure Agent.

Restriction: These steps do not work for Database Ingestion and Replication.

Contact your network administrator for the proxy settings.

1. Navigate to the following directory:

```
<Secure Agent installation directory>/apps/agentcore/conf
```

2. To update the `proxy.ini` file, add the following parameters and specify appropriate values for each parameter:

```
InfaAgent.ProxyHost=<proxy_server_hostname>
InfaAgent.ProxyPort=<proxy_server_port>
InfaAgent.ProxyUser=<user_name>
InfaAgent.ProxyPassword=<password>
InfaAgent.ProxyPasswordEncrypted=false
```

For example,

```
InfaAgent.ProxyHost=INW2PF0MT01V
InfaAgent.ProxyPort=808
InfaAgent.ProxyUser=user06
InfaAgent.ProxyPassword=user06
InfaAgent.ProxyPasswordEncrypted=false
```

3. Log in to Informatica Intelligent Cloud Services.
4. Open Administrator and select **Runtime Environments**.
5. Select the Secure Agent for which you want to configure a proxy server.
6. On the upper-right corner of the page, click **Edit**.
7. In the **System Configuration Details** section, select the **Type** as **Agent** for the CMI Streaming Agent Service.
8. To use a proxy server, add the following parameters in any **JVMOption** field and specify appropriate values for each parameter:

Parameter	Description
-Dproxy.host=	Host name of the outgoing HTTPS proxy server.
-Dproxy.port=	Port number of the outgoing HTTPS proxy server.
-Dproxy.user=	User name for the HTTPS proxy server.
-Dproxy.password=	Password for the HTTPS proxy server.

Note: You must specify the parameter and the value for the parameter enclosed in single quotation marks.

For example,

```
JVMOption1='-Dproxy.host=INPQ8583WI29'
JVMOption2='-Dproxy.port=8081'
JVMOption3='-Dproxy.user=adminuser'
JVMOption4='-Dproxy.password=password'
```

Note: You can configure only five **JVMOption** fields in the **System Configuration Details** section. To configure the remaining parameters, you must add the **JVMOption** fields in the **Custom Configuration Details** section. In the **Custom Configuration Details** section, select the **Type** as **Agent** for the CMI Streaming Agent Service, add the **JVMOption** fields, and specify the remaining parameters and appropriate values for each parameter.

9. Click **Save**.

The Secure Agent restarts to apply the settings.

Note: The session log does not record the proxy server details even if you have configured a proxy server.

Google PubSub - Streaming Ingestion and Replication connection properties

When you define a Google PubSub Streaming Ingestion and Replication connection, you must configure connection properties. You can use this connection type in streaming ingestion and replication tasks, which you configure in the Data Ingestion and Replication service.

The following table describes the Google PubSub connection properties:

Property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~`!\$%^&*() - + = { } \ : ; " ' < , > . ? /
Description	Optional. Description that you use to identify the connection. The description must not exceed 4,000 characters.
Type	The Google PubSub connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Client Email	The <code>client_email</code> value available in the JSON file that you download after you create a service account.
Client ID	The <code>client_id</code> value available in the JSON file that you download after you create a service account.
Private Key ID	The <code>private_key_id</code> value available in the JSON file that you download after you create a service account.
Private Key	The <code>private_key</code> value available in the JSON file that you download after you create a service account.
Project ID	The <code>project_id</code> value available in the JSON file that you download after you create a service account.

Note: The test connection for the Google PubSub connector does not fail even if you enter incorrect values for **Client ID** and **Private Key ID**.

Hadoop Files V2 connection properties

When you set up a Hadoop Files V2 connection, you must configure the connection properties.

The following table describes the Hadoop Files V2 connection properties:

Connection property	Description
Connection Name	Name of the Hadoop Files V2 connection.
Description	Description of the connection. The description cannot exceed 765 characters.
Type	Type of connection. Select Hadoop Files V2 .
Runtime Environment	The name of the runtime environment where you want to run the tasks.
User Name	Required to read data from HDFS. Enter a user name that has access to the single-node HDFS location to read data from or write data to.
NameNode URI	<p>The URI to access HDFS.</p> <p>Use the following format to specify the name node URI in Cloudera, Amazon EMR, and Hortonworks distributions:</p> <pre>hdfs://<namenode>:<port>/</pre> <p>where,</p> <ul style="list-style-type: none">- <namenode> is the host name or IP address of the name node.- <port> is the port that the name node listens for remote procedure calls (RPC). <p>To connect to the Hadoop cluster, specify the name node port <code>fs.defaultFS</code>.</p> <p>If the Hadoop cluster is configured for high availability, you must copy the <code>fs.defaultFS</code> value in the <code>core-site.xml</code> file and append <code>/</code> to specify the name node URI.</p> <p>For example, the following snippet shows the <code>fs.defaultFS</code> value in a sample <code>core-site.xml</code> file:</p> <pre><property> <name>fs.defaultFS</name> <value>hdfs://nameservice1</value> <source>core-site.xml</source> </property></pre> <p>In the above snippet, the <code>fs.defaultFS</code> value is</p> <pre>hdfs://nameservice1</pre> <p>and the corresponding name node URI is</p> <pre>hdfs://nameservice1/</pre> <p>Note: Specify either the name node URI or the local path. Do not specify the name node URI if you want to read data from or write data to a local file system path.</p>

Connection property	Description
Local Path	<p>A local file system path to read and write data. Read the following conditions to specify the local path:</p> <ul style="list-style-type: none"> - You must enter NA in local path if you specify the name node URI. If the local path does not contain NA, the name node URI does not work. - If you specify the name node URI and local path, the local path takes the preference. The connection uses the local path to run all tasks. - If you leave the local path blank, the agent configures the root directory (/) in the connection. The connection uses the local path to run all tasks. - If the file or directory is in the local system, enter the fully qualified path of the file or directory. <p>For example, /user/testdir specifies the location of a directory in the local system.</p> <p>Default value for Local Path is NA.</p>
Configuration Files Path	<p>The directory that contains the Hadoop configuration files.</p> <p>Note: Copy the core-site.xml, hdfs-site.xml, and hive-site.xml from the Hadoop cluster and add them to a folder in Linux Box.</p>
Keytab File	The file that contains encrypted keys and Kerberos principals to authenticate the machine.
Principal Name	Users assigned to the superuser privilege can perform all the tasks that a user with the administrator privilege can perform.
Impersonation Username	You can enable different users to run mappings in a Hadoop cluster that uses Kerberos authentication or connect to sources and targets that use Kerberos authentication. To enable different users to run mappings or connect to big data sources and targets, you must configure user impersonation.

Note: When you read from or write to remote files, the **NameNode URI** and **Configuration Files Path** fields are mandatory. When you read from or write to local files, you require only the **Local Path** field.

JDBC V2 connection properties

When you set up a JDBC V2 connection, configure the connection properties.

The following table describes the JDBC V2 connection properties:

Property	Description
Connection Name	Name of the connection.
Description	Description of the connection.
Type	Type of connection. Select JDBC V2 from the list.
Runtime Environment	The name of the runtime environment where you want to run tasks. Select a Secure Agent.
User Name	The user name to connect to the database.

Property	Description
Password	The password for the database user name.
Schema Name	Optional. The schema name. If you don't specify the schema name, all the schemas available in the database are listed. To read from or write to Oracle public synonyms, enter PUBLIC.
JDBC Driver Class Name	Name of the JDBC driver class. To connect to Aurora PostgreSQL, specify the following driver class name: org.postgresql.Driver For more information about which driver class to use with specific databases, see the corresponding third-party vendor documentation.
Connection String	Connection string to connect to the database. Use the following format to specify the connection string: <code>jdbc:<subprotocol>:<subname></code> For example, the connection string for the Aurora PostgreSQL database type is <code>jdbc:postgresql://<host>:<port>[/dbname]</code> . For more information about the connection string to use with specific drivers, see the corresponding third-party vendor documentation.
Additional Security Properties	Masks sensitive and confidential data of the connection string that you don't want to display in the session log. Specify the part of the connection string that you want to mask. When you create a connection, the string you enter in this field appends to the string that you specified in the Connection String field.
Database Type	The database type to which you want to connect. You can select one of the following database types: <ul style="list-style-type: none"> - PostgreSQL. Connect to the Aurora PostgreSQL database hosted in the Amazon Web Services or the Microsoft Azure environment. - Azure SQL Database. Connect to Azure SQL Database hosted in the Microsoft Azure environment. - Others. Connect to any database that supports the Type 4 JDBC driver. Default is Others.
Enable Auto Commit ¹	Specifies whether the driver supports connections to automatically commit data to the database when you run an SQL statement. When disabled, the driver does not support connections to automatically commit data even if the auto-commit mode is enabled in the JDBC driver. Default is disabled.
Support Mixed-Case Identifiers	Indicates whether the database supports case-sensitive identifiers. When enabled, the Secure Agent encloses all identifiers within the character selected for the SQL Identifier Character property. Default is disabled.

Property	Description
SQL Identifier Character	<p>Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type.</p> <p>Select None if the database uses regular identifiers. When the Secure Agent generates SQL queries, it does not place delimited characters around any identifiers.</p> <p>Select a character if the database uses delimited identifiers. When the Secure Agent generates SQL queries, it encloses delimited identifiers within this character.</p>
¹ Doesn't apply to mappings in advanced mode.	

JMS connection properties

When you set up a JMS connection, you must configure the connection properties.

The following table describes the connection properties for the JMS connection:

Property	Description
Connection Name	<p>Name of the connection.</p> <p>The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:</p> <p>~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /</p>
Description	<p>Optional. Description that you can use to identity the connection.</p> <p>The description cannot exceed 4,000 characters.</p>
Type	<p>The JMS connection type.</p> <p>If you do not see the connection type, go to the Add-On Connectors page to install the connector.</p>
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Connection URL	<p>URL of the JNDI naming provider.</p> <p>For example, in IBM MQ it is the directory location that contains the .bindings file.</p>
JNDI User Name	Optional. User name to connect to the JNDI context factory.
JNDI Password	Optional. The password of the user account that you use to connect to the JNDI context factory.
JNDI Context Factory	<p>The JMS provider specific initial JNDI context factory implementation for connecting to the JNDI service. This value is a fully qualified class name of the Initial Context Factory.</p> <p>For example, the class name of the Initial Context Factory for ActiveMQ is</p> <pre>org.apache.activemq.jndi.ActiveMQInitialContextFactory</pre> <p>For more information, see the documentation of the JMS provider.</p>

Property	Description
JNDI Package Prefixes	A colon-delimited list of package prefixes to use when loading URL context factories. These are the package prefixes for the name of the factory class that will create a URL context factory. For more information about the values, see the documentation of the JMS provider.
JMS Connection Factory	The name of the object in the JNDI server that enables the JMS Client to create JMS connections. For example, <code>jms/QCF</code> or <code>jmsSalesSystem</code> .
JMS Connection User Name	Optional. User name to connect to the JMS connection factory.
JMS Connection Password	Optional. The password of the user account that you use to connect to the JMS connection factory.

Note: Ensure to copy the external JMS JAR files to the following location:

```
<Secure_Agent_home>/ext/connectors/thirdparty/infa.jms
```

After copying the external JMS JAR files, restart the Secure Agent.

Kafka connection properties

When you set up a Kafka connection, configure the connection properties.

The following table describes the Kafka connection properties:

Property	Description
Connection Name	Name of the connection. The name is not case sensitive. It must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
Description	Optional. Description that you use to identity the connection. The description cannot exceed 4,000 characters.
Type	The Kafka connection type. If you do not see the connection type, go to the Add-On Connectors page in Administrator to install the connector.
Runtime Environment	Name of the runtime environment where you want to run tasks.

Property	Description
Kafka Broker List	<p>Comma-separated list of the Kafka brokers.</p> <p>To list a Kafka broker, use the following format:</p> <pre><HostName>:<PortNumber></pre> <p>Note: When you connect to a Kafka broker over SSL, you must specify the fully qualified domain name for the host name. Otherwise, the test connection fails with SSL handshake error.</p>
Retry Timeout	<p>Optional. Number of seconds after which the Secure Agent attempts to reconnect to the Kafka broker to read or write data.</p> <p>Default is 180 seconds.</p> <p>This property is not used by Database Ingestion and Replication. You can specify an equivalent Kafka property in Additional Connection Properties.</p>
Kafka Broker Version	<p>Kafka message broker version. The only valid value is Apache 0.10.1.1 and above.</p> <p>Optional for a streaming ingestion and replication task.</p>
Additional Connection Properties	<p>Optional. Comma-separated list of additional configuration properties of the Kafka producer or consumer.</p> <p>For a streaming ingestion and replication task, ensure that you set the <code><kerberos name></code> property if you configure <code><Security Protocol></code> as <code>SASL_PLAINTEXT</code> or <code>SASL_SSL</code>.</p> <p>For a database ingestion and replication task, if you want to specify a security protocol and properties, specify them here instead of in the Additional Security Properties property. For example: <code>security.protocol=SSL,ssl.truststore.location=/opt/kafka/config/kafka.truststore.jks,ssl.truststore.password=<trustore_password></code>.</p>
Schema Registry URL	<p>Location and port of the Confluent schema registry service to access Avro sources and targets in Kafka.</p> <p>To list a schema registry URL, use the following format:</p> <pre><https>://<HostName or IP>:<PortNumber></pre> <p>or</p> <pre><http>://<HostName or IP>:<PortNumber></pre> <p>Example for the schema registry URL:</p> <pre>https://kafkarnd.informatika.com:8082</pre> <p>or</p> <pre>http://10.65.146.181:8084</pre> <p>Applies only when you import a Kafka topic in Avro format that uses the Confluent schema registry to store the metadata.</p> <p>This property is not used by Database Ingestion and Replication. You can specify an equivalent Kafka property in Additional Connection Properties.</p>
SSL Mode	<p>Required. Determines the encryption type to use for the connection.</p> <p>You can choose a mode from the following SSL modes:</p> <ul style="list-style-type: none"> - Disabled. Establishes an unencrypted connection to the Kafka broker. - One-way. Establishes an encrypted connection to the Kafka broker using truststore file and truststore password. - Two-way. Establishes an encrypted connection to the Kafka broker using truststore file, truststore password, keystore file, and keystore password. <p>This property is not used by Database Ingestion and Replication. You can specify an equivalent Kafka property in Additional Connection Properties.</p>

Property	Description
SSL TrustStore File Path	Required when you use the one-way or two-way SSL mode. Absolute path and file name of the SSL truststore file that contains the SSL certificate to connect to the Kafka broker.
SSL TrustStore Password	Required when you use the one-way or two-way SSL mode. Password for the SSL truststore.
SSL KeyStore File Path	Required when you use the two-way SSL mode. Absolute path and file name of the SSL keystore file that contains private keys and certificates to connect to the Kafka broker.
SSL KeyStore Password	Required when you use the two-way SSL mode. Password for the SSL keystore.
Additional Security Properties	Optional. Comma-separated list of additional configuration properties to connect to the Kafka broker in a secure way. If you specify two different values for the same property in Additional Connection Properties and Additional Security Properties , the value in Additional Security Properties overrides the value in Additional Connection Properties . This property is not used by Database Ingestion and Replication. You can specify a security protocol and properties in Additional Connection Properties .

Schema Registry Security Configuration Properties

When you configure the **Schema Registry URL** connection property, you can configure the schema registry security configuration properties. These properties apply only to mappings in advanced mode. You can configure one-way SSL, two-way SSL, and basic authentication to connect to the Confluent schema registry in a secure way.

The following table describes the security properties for the Kafka connection when you use the Confluent schema registry:

Property	Description
SSL Mode Schema Registry ¹	Required. Determines the encryption type to use for the connection. You can choose a mode from the following SSL modes: <ul style="list-style-type: none"> - Disabled. Establishes an unencrypted connection to the Confluent schema registry. - One-way. Establishes an encrypted connection to the Confluent schema registry using truststore file and truststore password. - Two-way. Establishes an encrypted connection to the Confluent schema registry using truststore file, truststore password, keystore file, and keystore password. This property is not used by Database Ingestion and Replication. You can specify an equivalent Kafka property in Additional Connection Properties .
SSL TrustStore File Path Schema Registry ¹	Required when you use the one-way or two-way SSL mode. Absolute path and file name of the SSL truststore file that contains the SSL certificate to connect to the Confluent schema registry.

Property	Description
SSL TrustStore Password Schema Registry ¹	Required when you use the one-way or two-way SSL mode. Password for the SSL truststore.
SSL KeyStore File Path Schema Registry ¹	Required when you use the two-way SSL mode. Absolute path and file name of the SSL keystore file that contains private keys and certificates to connect to the Confluent schema registry.
SSL KeyStore Password Schema Registry ¹	Required when you use the two-way SSL mode. Password for the SSL keystore.
Additional Security Properties Schema Registry ²	Optional. Comma-separated list of additional security properties to connect to the Confluent schema registry in a secure way. For example, when you configure basic authentication to establish a secure communication with Confluent schema registry, specify the following value: <code>basic.auth.credentials.source=USER_INFO,basic.auth.user.info=<username>:<password></code> If you specify two different values for the same property in Additional Connection Properties and Additional Security Properties Schema Registry , the value in Additional Security Properties Schema Registry overrides the value in Additional Connection Properties . This property is not used by Database Ingestion and Replication.
¹ Applies only to mappings in advanced mode. ² Applies to both mappings and mappings in advanced mode.	

Configuring the krb5.conf file to read data from or write to a Kerberised Kafka cluster

To read from or write to a Kerberised Kafka cluster, configure the default realm, KDC, and Kafka advanced source or target properties.

You can configure Kerberos authentication for a Kafka client by placing the required Kerberos configuration files on the Secure Agent machine and specifying the required JAAS configuration in the Kafka connection. The JAAS configuration defines the keytab and principal details that the Kafka broker must use to authenticate the Kafka client.

Note: This topic is not applicable to Application Ingestion and Replication and Database Ingestion and Replication. Application Ingestion and Replication and Database Ingestion and Replication do not yet support this functionality.

Before you read from or write to a Kerberised Kafka cluster, perform the following tasks:

1. Ensure that you have the `krb5.conf` file for the Kerberised Kafka cluster.
2. Configure the default realm and KDC. If the default `/etc/krb5.conf` file is not configured or you want to change the configuration, add the following lines to the `/etc/krb5.conf` file:

```
[libdefaults]
default_realm = <REALM NAME>
```

```

dns_lookup_realm = false
dns_lookup_kdc = false
ticket_lifetime = 24h
renew_lifetime = 7d
forwardable = true

[realms]
<REALM NAME> = {
kdc = <Location where KDC is installed>
admin_server = <Location where KDC is installed>
}

[domain_realm]
.<domain name or hostname> = <KERBEROS DOMAIN NAME>
<domain name or hostname> = <KERBEROS DOMAIN NAME>

```

3. To pass a static JAAS configuration file into the JVM using the `java.security.auth.login.config` property at runtime, perform the following tasks:

- a. Ensure that you have JAAS configuration file.

For information about creating JAAS configuration and configuring keytab for Kafka clients, see the Apache Kafka documentation at <https://kafka.apache.org/0101/documentation/#security>

For example, the JAAS configuration file can contain the following lines of configuration:

```

//Kafka Client Authentication. Used for client to kafka broker connection
KafkaClient {
com.sun.security.auth.module.Krb5LoginModule required
doNotPrompt=true
useKeyTab=true
storeKey=true
keyTab="<path to Kafka keytab file>/<Kafka keytab file name>"
principal="<principal name>"
client=true
};

```

- b. Place the JAAS config file and keytab file in the same location on all the secure agents.

Informatica recommends that you place the files in a location that is accessible by all the secure agents in the runtime environment. For example, `/etc` or `/temp`.

- c. Configure the following properties:

Kafka connection

Configure the **Additional Connection Properties** in a Kafka connection and specify the value in the following format:

```

security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=kafka,sasl.mechanism=GSSAPI

```

Sources

Configure the **Consumer Configuration Properties** in the advanced source properties to override the value specified in **Additional Connection Properties** in a Kafka connection. Specify the value in the following format:

```

security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=kafka,sasl.mechanism=GSSAPI

```

Targets

Configure the **Producer Configuration Properties** in the advanced target properties to override the value specified in **Additional Connection Properties** in a Kafka connection. Specify the value in the following format:

```

security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=kafka,sasl.mechanism=GSSAPI

```

4. To embed the JAAS configuration in the `sasl.jaas.config` configuration property, configure the following properties:

Kafka connection

Configure the **Additional Connection Properties** in a Kafka connection and specify the value in the following format:

```
security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=kafka,sasl.mechanism=
GSSAPI,
sasl.jaas.config=com.sun.security.auth.module.Krb5LoginModule required
useKeyTab=true
storeKey=true doNotPrompt=true serviceName="<service_name>" keyTab="<location of
keytab file>"
client=true principal="<principal_name>;
```

Sources

Configure the **Consumer Configuration Properties** in the advanced source properties to override the value specified in **Kerberos Configuration Properties** in a Kafka connection. Specify the value in the following format:

```
security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=kafka,sasl.mechanism=
GSSAPI,
sasl.jaas.config=com.sun.security.auth.module.Krb5LoginModule required
useKeyTab=true
storeKey=true doNotPrompt=true serviceName="<service_name>" keyTab="<location of
keytab file>"
client=true principal="<principal_name>;
```

Targets

Configure the **Producer Configuration Properties** in the advanced target properties to override the value specified in **Kerberos Configuration Properties** in a Kafka connection. Specify the value in the following format:

```
security.protocol=SASL_PLAINTEXT,sasl.kerberos.service.name=kafka,sasl.mechanism=
GSSAPI,
sasl.jaas.config=com.sun.security.auth.module.Krb5LoginModule required
useKeyTab=true
storeKey=true doNotPrompt=true serviceName="<service_name>" keyTab="<location of
keytab file>"
client=true principal="<principal_name>;
```

Configuring SASL PLAIN authentication for a Kafka cluster

In the Kafka connection, you can configure PLAIN security for the Kafka broker to connect to a Kafka broker. To read data from or write data to a Kafka broker with SASL PLAIN authentication, configure the Kafka connection properties. To override the properties defined in the Kafka connection, you can configure the advanced source or target properties.

You can configure SASL PLAIN authentication so that the Kafka broker can authenticate the Kafka producer and the Kafka consumer. Kafka uses the Java Authentication and Authorization Service (JAAS) for SASL PLAIN authentication. To enable SASL PLAIN authentication, you must specify the SASL mechanism as PLAIN. You must also provide the formatted JAAS configuration that the Kafka broker must use for authentication. The JAAS configuration defines the username, password, that the Kafka broker must use to authenticate the Kafka client.

Note: This topic is not applicable to Application Ingestion and Replication and Database Ingestion and Replication. Application Ingestion and Replication and Database Ingestion and Replication do not yet support this functionality.

Configure the following properties:

Kafka connection

Configure the **Additional Connection Properties** or **Additional Security Properties** property in the Kafka connection and specify the value in the following format:

```
security.protocol=SASL_SSL,sasl.mechanism=PLAIN,sasl.jaas.config=org.apache.kafka.com  
mon.security.plain.PlainLoginModule required username="<username>"  
password="<password>"
```

In the **Security Configuration Section**, select **One-Way** as the **SSL Mode** and specify the SSL TrustStore File Path and SSL TrustStore Password.

Sources

Configure the **Consumer Configuration Properties** property in the advanced source properties to override the value that you specified in the **Additional Connection Properties** property in the Kafka connection. Specify the value in the following format:

```
security.protocol=SASL_SSL,sasl.mechanism=PLAIN,sasl.jaas.config=org.apache.kafka.com  
mon.security.plain.PlainLoginModule required username="<username>"  
password="<password>"
```

Targets

Configure the **Producer Configuration Properties** property in the advanced target properties to override the value that you specified in the **Additional Connection Properties** property in the Kafka connection. Specify the value in the following format:

```
security.protocol=SASL_SSL,sasl.mechanism=PLAIN,sasl.jaas.config=org.apache.kafka.com  
mon.security.plain.PlainLoginModule required username="<username>"  
password="<password>"
```

Configuring SASL PLAIN authentication for an Azure Event Hub Kafka broker

In the Kafka connection, you can configure PLAIN security for the Kafka broker to connect to an Azure Event Hub Kafka broker. When you connect to an Azure Event Hub Kafka broker, the password defines the endpoint URL that contains the fully qualified domain name (FQDN) of the Event Hub namespace, shared access key name, and shared access key required to connect to an Azure Event Hub Kafka broker. Configure the SSL Mode as One-Way and provide the path to a trusted root certificate on your file system for SSL TrustStore File Path.

To connect to an Azure Event Hub Kafka broker, configure any of the above properties and specify the value in the following format:

```
security.protocol=SASL_SSL,sasl.mechanism=PLAIN,sasl.kerberos.service.name=Kafka,sasl.jaa  
s.config=org.apache.kafka.common.security.plain.PlainLoginModule required  
username="$ConnectionString" password="Endpoint=sb://<FQDN>/;SharedAccessKeyName=<key  
name>;SharedAccessKey=<shared access key>"
```

Configuring SASL_SSL authentication for a Cloud Confluent Kafka cluster

In the Kafka connection, you can configure SSL security for encryption and authentication while connecting to a Kafka broker. To read data from or write data to a Confluent Kafka broker with SASL_SSL authentication, configure the Kafka connection properties. To override the properties defined in the Kafka connection, you can configure the advanced source or target properties.

Note: This topic is not applicable to Application Ingestion and Replication and Database Ingestion and Replication. Application Ingestion and Replication and Database Ingestion and Replication do not yet support this functionality.

Configure the following properties:

Property	Values
Additional Connection Properties	<code>security.protocol=SASL_SSL,sasl.kerberos.service.name=kafka,ssl.endpoint.identification.algorithm=required username=<> password=<></code>
SSL Mode	One-way
SSL TrustStore File Path	Use <code>cacert</code> file of agent JDK. For example: <code>/root/staging/infaagent/jdk/jre/lib/security/cacerts</code>
SSL TrustStore Password	Password for the SSL truststore.

Connecting to Amazon Managed Streaming for Apache Kafka

In the Kafka connection, you can configure PLAINTEXT or TLS encryption to connect to an Amazon Managed Streaming for Apache Kafka broker. To read data from or write data to an Amazon Managed Streaming for Apache Kafka broker, configure the Kafka connection properties.

Configure the **Kafka Broker List** property in the Kafka connection and specify the comma-separated list of Kafka brokers that you want to connect to in the following format:

`<HostName>:<PortNumber>`

Configure TLS encryption to securely connect the Kafka broker to the Kafka producer and the Kafka consumer. To configure TLS encryption for an Amazon Managed Streaming for Apache Kafka broker, configure the following properties:

Property	Values
Additional Connection Properties	<code>security.protocol=SSL</code>
SSL Mode	One-way or Two-way.
SSL TrustStore File Path	Required when you use the one-way or two-way SSL mode. Absolute path and file name of the SSL truststore file.
SSL TrustStore Password	Required when you use the one-way or two-way SSL mode. Password for the SSL truststore.
SSL KeyStore File Path	Required when you use the two-way SSL mode. Absolute path and file name of the SSL keystore file that contains private keys and certificates that the Kafka broker validates against the Kafka cluster certificate.
SSL KeyStore Password	Required when you use the two-way SSL mode. Password for the SSL keystore.

When you run a mapping that runs on an advanced cluster and connect to an Amazon Managed Streaming for Apache Kafka broker, configure the Kafka broker using SASL_SSL authentication with Salted Challenge

Response Authentication Mechanism (SCRAM). To read data from or write data to an Amazon Managed Streaming for Apache Kafka broker with SASL_SSL authentication, configure the following properties:

Property	Values
Additional Connection Properties	<code>security.protocol=SASL_SSL,sasl.mechanism=SCRAM-SHA-512,sasl.jaas.config=org.apache.kafka.common.security.scram.ScramLoginModule required username="<username>" password="<password>";</code>
SSL Mode	One-way or Two-way.
SSL TrustStore File Path	Required when you use the one-way or two-way SSL mode. Absolute path and file name of the SSL truststore file.
SSL TrustStore Password	Required when you use the one-way or two-way SSL mode. Password for the SSL truststore.
SSL KeyStore File Path	Required when you use the two-way SSL mode. Absolute path and file name of the SSL keystore file that contains private keys and certificates that the Kafka broker validates against the Kafka cluster certificate.
SSL KeyStore Password	Required when you use the two-way SSL mode. Password for the SSL keystore.

Marketo V3 connection properties

When you set up a Marketo V3 connection, configure the connection properties.

The following table describes the Marketo V3 connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Marketo V3 connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
client_ID	The client ID of the custom service required to generate a valid access token.
client_secret	The client secret of the Marketo custom service required to generate a valid access token.
grant_type	Marketo supports only the client_credentials grant type.

Property	Description
REST API URL	The URL has the following format: https://<Host name of the Marketo Rest API Server>. Contact the Marketo Administrator for the REST API URL.
Bypass Proxy	Note: This property is not applicable to connections configured for application ingestion and replication tasks.

Microsoft Azure Blob Storage V3 connection properties

When you set up a Microsoft Azure Blob Storage V3 connection, configure the connection properties.

The following table describes the Microsoft Azure Blob Storage V3 connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Microsoft Azure Blob Storage V3 connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. Specify a Secure Agent, Hosted Agent, or serverless runtime environment.
Account Name	Microsoft Azure Blob Storage account name.
Authentication Type	Authentication type to access the Microsoft Azure Blob Storage account. Select one of the following options: <ul style="list-style-type: none"> - Shared Key Authentication. Uses the account key to connect to Microsoft Azure Blob Storage. - Shared Access Signature. Uses the SAS token to connect to Microsoft Azure Blob Storage. Use the SAS token to grant access to the resources in the storage account or container for a specific time range without sharing the account key. Note: The file ingestion and replication task fails if this option is on a container level and if you use a different container.
Account Key	Applies to shared key authentication. The account key for the Microsoft Azure Blob Storage account.
SAS Token	Applies to shared access signature. The shared access signature token generated in the Azure portal.

Property	Description
Container Name	Microsoft Azure Blob Storage container name.
Endpoint Suffix	Type of Microsoft Azure endpoints. Select one of the following options: <ul style="list-style-type: none"> - core.windows.net. Connects to Azure endpoints. - core.usgovcloudapi.net. Connects to Azure Government endpoints. - core.chinacloudapi.cn. Not applicable. Default is core.windows.net.

Microsoft Azure Data Lake Storage Gen2 connection properties

When you set up a Microsoft Azure Data Lake Storage Gen2 connection, configure the connection properties.

The following table describes the Microsoft Azure Data Lake Storage Gen2 connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ - + , Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Microsoft Azure Data Lake Storage Gen2 connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run a database ingestion and replication task or streaming ingestion and replication task on a Hosted Agent or serverless runtime environment.
Account Name	Microsoft Azure Data Lake Storage Gen2 account name or the service name.
Authentication Type	Authentication type to access the Microsoft Azure Data Lake Storage Gen2 account. Select one of the following options: <ul style="list-style-type: none"> - Service Principal Authentication. Uses the client ID, client secret, and tenant ID to connect to Microsoft Azure Data Lake Storage Gen2. - Shared Key Authentication. Uses the account key to connect to Microsoft Azure Data Lake Storage Gen2. - Managed Identity Authentication. Select to authenticate using identities that are assigned to applications in Azure to access Azure resources in Microsoft Azure Data Lake Storage Gen2. Note: Streaming Ingestion and Replication does not support shared key authentication or managed identity authentication.

Property	Description
Client ID	<p>Applies to Service Principal Authentication and Managed Identity Authentication.</p> <p>The client ID of your application.</p> <p>To use service principal authentication, specify the application ID or client ID for your application registered in the Azure Active Directory.</p> <p>To use managed identity authentication, specify the client ID for the user-assigned managed identity. If the permission is provided by system-assigned managed identity, leave the field empty. If there is no system-assigned identity but only a single user-assigned managed identity, you may also leave the field empty.</p>
Client Secret	<p>Applies to Service Principal Authentication.</p> <p>The client secret key to complete the OAuth authentication in the Azure Active Directory.</p>
Tenant ID	<p>Applies to Service Principal Authentication.</p> <p>The directory ID of the Azure Active Directory.</p>
Account Key	<p>Applies to Shared Key Authentication.</p> <p>The account key for the Microsoft Azure Data Lake Storage Gen2 account.</p>
File System Name	<p>The name of the file system in the Microsoft Azure Data Lake Storage Gen2 account.</p>
Directory Path	<p>The path of an existing directory without the file system name.</p> <p>You can select one of the following syntax:</p> <ul style="list-style-type: none"> - / for root directory - /dir1 - dir1/dir2 <p>There is no default directory.</p>
Adls Gen2 End-point	<p>The type of Microsoft Azure endpoints.</p> <p>Select one of the following endpoints:</p> <ul style="list-style-type: none"> - core.windows.net. Connects to Azure endpoints. - core.usgovcloudapi.net. Connects to US government Microsoft Azure Data Lake storage Gen2 endpoints. - core.chinacloudapi.cn. Connects to Microsoft Azure Data Lake storage Gen2 endpoints in the China region. <p>Default is core.windows.net.</p>

Microsoft Azure Event Hub connection properties

When you set up an Azure Event Hub connection, you must configure the connection properties.

The following table describes the Azure Event Hub connection properties:

Property	Description
Connection Name	Name of the connection. The name is not case sensitive. It must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [}] \ : ; " ' < , > . ? /
Description	Optional. Description that you use to identify the connection. The description cannot exceed 4,000 characters.
Type	The Azure Event Hub connection type. If you do not see the connection type, go to the Add-On Connectors page in Administrator to install the connector.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Tenant ID	The ID of the tenant that the data belongs to. This ID is the Directory ID of the Azure Active Directory.
Subscription ID	The ID of the Azure subscription.
Resource Group Name	The name of the Azure resource group associated with the Event Hub namespace.
Client Application ID	The ID of the application created under the Azure Active Directory.
Client Secret Key	The secret key generated for the application.
Event Hub Namespace	The name of the Event Hub namespace that is associated with the resource group name.
Shared Access Policy Name	Optional. The name of the Event Hub Namespace Shared Access Policy. The policy must apply to all data objects that are associated with this connection. To read from Event Hubs, you must have Listen permission. To write to an Event Hub, the policy must have Send permission.
Shared Access Policy Primary Key	Optional. The primary key of the Event Hub Namespace Shared Access Policy.

Microsoft Azure Synapse Analytics Database Ingestion connection properties

When you define a Microsoft Azure Synapse Analytics Database Ingestion connection, you must configure connection properties. You can use this connection type in application ingestion and replication tasks and database ingestion and replication tasks, which you configure in the Data Ingestion and Replication service.

Note: Some properties are for Microsoft Azure Data Lake Storage Gen2. Application Ingestion and Replication and Database Ingestion and Replication use Microsoft Azure Data Lake Storage Gen2 to stage data in files before sending the data to the Microsoft Azure Synapse Analytics target tables.

The following table describes the connection properties:

Property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Ensure that the type is Microsoft Azure Synapse Analytics Database Ingestion .
Runtime Environment	The name of the runtime environment where you want to run the application ingestion and replication tasks or database ingestion and replication tasks. You define runtime environments in Administrator. Note: You cannot run application ingestion and replication tasks or database ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Azure Synapse Analytics JDBC URL	The Microsoft Azure Synapse Analytics (formerly SQL Data Warehouse) JDBC connection string. Example connection string for Microsoft SQL Server authentication: <code>jdbc:sqlserver://server.database.windows.net:1433;database=database</code> Example connection string for Azure Active Directory (AAD) authentication: <code>jdbc:sqlserver://server.database.windows.net:1433;database=database;encrypt=true;trustServerCertificate=false;hostNameInCertificate=*.database.windows.net;loginTimeout=30;Authentication=ActiveDirectoryPassword;</code> Note: The default authentication type is Microsoft SQL Server authentication.
Azure Synapse Analytics JDBC Username	The user name to use for connecting to the Microsoft Azure Synapse Analytics account. Provide the AAD user name for AAD authentication.
Azure Synapse Analytics JDBC Password	The password to use for connecting to the Microsoft Azure Synapse Analytics account.
Azure Synapse Analytics Schema Name	The name of the schema in the Microsoft Azure Synapse Analytics target.

Property	Description
ADLS Gen2 Account Name	The name of the Microsoft Azure Data Lake Storage Gen2 account.
Client Id	The ID of your client application for completing the OAuth Authentication in the Active Directory.
Client Secret	The client secret key for completing the OAuth Authentication in the Active Directory.
Directory	The Microsoft Azure Data Lake Storage Gen2 directory that Application Ingestion and Replication and Database Ingestion and Replication use to stage data in files. The default is the root directory.
Filesystem Name	The name of an existing file system in the Microsoft Azure Data Lake Storage Gen2 account.
Tenant ID	The Directory ID of the Azure Active Directory.

Microsoft Azure Synapse SQL connection properties

When you set up a Microsoft Azure Synapse SQL connection, configure the connection properties.

The following table describes the Microsoft Azure Synapse SQL connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Microsoft Azure Synapse SQL connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. For application, database, and streaming ingestion and replication tasks, you cannot use the Hosted Agent or serverless runtime environments. For file ingestion and replication tasks, you can use the Hosted Agent but not serverless environments.

Property	Description
Azure DW JDBC URL	<p>The Microsoft Azure Synapse SQL JDBC connection string.</p> <p>Enter the connection string in the following format for Microsoft SQL Server authentication:</p> <pre>jdbc:sqlserver://<Server>.database.windows.net:1433; database=<Database></pre> <p>Enter the connection string in the following format for Azure Active Directory (AAD) authentication:</p> <pre>jdbc:sqlserver://<Server>.database.windows.net:1433; database=<Database>;encrypt=true;trustServerCertificate=false; hostNameInCertificate=*.database.windows.net;loginTimeout=30; Authentication=ActiveDirectoryPassword;</pre> <p>Enter the connection string in the following format for Managed Identity authentication:</p> <pre>jdbc:sqlserver://<Server>.database.windows.net:1433; database=<Database>;Authentication=ActiveDirectoryMsi;</pre> <p>Default is Microsoft SQL Server authentication.</p>
Azure DW JDBC Username	User name to connect to the Microsoft Azure Synapse SQL account. Provide AAD user name for AAD authentication.
Azure DW JDBC Password	Password to connect to the Microsoft Azure Synapse SQL account. For AAD authentication, provide the password of AAD user.
Azure DW Schema Name	Name of the schema in Microsoft Azure Synapse SQL.
Azure DW Client ID	<p>Required if you want to use the user-assigned managed identity for Managed Identity Authentication to connect to Microsoft Azure Synapse SQL.</p> <p>The client ID of the user-assigned managed identity.</p> <p>If the managed identity is system-assigned, leave the field empty.</p>
Azure Storage Type	<p>Type of Azure storage to stage the files.</p> <p>Select one of the following storage types:</p> <ul style="list-style-type: none"> - Azure Blob. Uses Microsoft Azure Blob Storage to stage the files. - ADLS Gen2. Uses Microsoft Azure Data Lake Storage Gen2 to stage the files. <p>Default is Azure Blob.</p>
Authentication Type	<p>Authentication type to connect to Azure storage to stage the files.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> - Shared Key Authentication. Uses the account name and account key to connect to Microsoft Azure Blob Storage or Microsoft Azure Data Lake Storage Gen2. - Service Principal Authentication. Applies to Microsoft Azure Data Lake Storage Gen2. Uses the client ID, client secret, and tenant ID to connect to Microsoft Azure Data Lake Storage Gen2. To use Service Principal authentication, register an application in the Azure Active Directory, generate a client secret, and then assign the Storage Blob Contributor role to the application. - Managed Identity Authentication. Applies to Microsoft Azure Data Lake Storage Gen2. Select to authenticate using identities that are assigned to applications in Azure to access Azure resources in Microsoft Azure Data Lake Storage Gen2. <p>In a file ingestion and replication task, if you select Microsoft Azure Synapse SQL with Managed Identity authentication type as the target, then you must select Microsoft Azure Data Lake Storage Gen2 as the source.</p>
Azure Blob Account Name	<p>Applies to Shared Key Authentication for Microsoft Azure Blob Storage.</p> <p>Name of the Microsoft Azure Blob Storage account to stage the files.</p>
Azure Blob Account Key	<p>Applies to Shared Key Authentication for Microsoft Azure Blob Storage.</p> <p>The Microsoft Azure Blob Storage access key to stage the files.</p>

Property	Description
Container Name	Applies to Microsoft Azure Blob Storage. The name of the container in the Azure Blob Storage account.
ADLS Gen2 Storage Account Name	Applies to Shared Key Authentication and Service Principal Authentication for Microsoft Azure Data Lake Storage Gen2. Name of the Microsoft Azure Data Lake Storage Gen2 account to stage the files.
ADLS Gen2 Account Key	Applies to Shared Key Authentication for Microsoft Azure Data Lake Storage Gen2. The Microsoft Azure Data Lake Storage Gen2 access key to stage the files.
Client ID	Applies to Service Principal Authentication and Managed Identity Authentication for Microsoft Azure Data Lake Storage Gen2. The client ID of your application. To use service principal authentication, enter the application ID or client ID for your application registered in the Azure Active Directory. To use managed identity authentication, enter the client ID for the user-assigned managed identity. If the managed identity is system-assigned, leave the field empty.
Client Secret	Applies to Service Principal Authentication for Microsoft Azure Data Lake Storage Gen2. The client secret for your application.
Tenant ID	Applies to Service Principal Authentication for Microsoft Azure Data Lake Storage Gen2. The directory ID or tenant ID for your application.
File System Name	Applies to Microsoft Azure Data Lake Storage Gen2. The name of the file system in the Microsoft Azure Data Lake Storage Gen2 account.
Blob End-point	Type of Microsoft Azure endpoints. Select one of the following endpoints: <ul style="list-style-type: none"> - core.windows.net. Connects to Azure endpoints. - core.usgovcloudapi.net. Connects to US Government Microsoft Azure Synapse SQL endpoints. - core.chinacloudapi.cn. Connects to Microsoft Azure Synapse SQL endpoints in the China region. Default is core.windows.net.
VNet Rule	Enable to connect to a Microsoft Azure Synapse SQL endpoint residing in a virtual network (VNet). When you use a serverless runtime environment, you cannot connect to a Microsoft Azure Synapse SQL endpoint residing in a virtual network.

Microsoft Dynamics 365 Mass Ingestion connection properties

When you set up a Microsoft Dynamics 365 Mass Ingestion connection, you must configure the connection properties.

The Microsoft Dynamics 365 Mass Ingestion connection requires a native application that is registered in Azure Active Directory (Azure AD) to access the Microsoft Dynamics 365 data. Before you configure the connection, you must register an application in Azure AD to allow the connection to access the Microsoft

Dynamics 365 data. For more information about registering an application in Azure AD, see the [Microsoft documentation](#).

The properties of a Microsoft Dynamics 365 Mass Ingestion connection vary based on the authentication method that you specify for the connection. When you create a connection, you can select one of the following authentication methods:

- **OAuth 2.0 Username-Password Flow:** Authenticates the connection by using the Microsoft Dynamics 365 account login credentials and the client ID of the application registered in Azure AD.
- **OAuth 2.0 Client Secret Flow:** Authenticates the connection by using the client ID and client secret of the application registered in Azure AD.
- **OAuth 2.0 JWT Bearer Flow:** Authenticates the connection by using a X509 Public Key Infrastructure (PKI) certificate and a JSON Web Token (JWT). Use this authentication method to gain secured access to Microsoft Dynamics 365 without sharing sensitive information, such as client secret and Microsoft Dynamics 365 account login credentials.

Connection properties for OAuth 2.0 Username-Password Flow authentication

The following table describes the connection properties for a Microsoft Dynamics 365 Mass Ingestion connection configured with OAuth 2.0 Username-Password Flow authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Select the Microsoft Dynamics 365 Mass Ingestion connection type.
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
User Name	User name of the Microsoft Dynamics 365 account.
Password	Password for the Microsoft Dynamics 365 account.
Client ID	Client ID of the application registered in Azure AD.
Resource URL	URL of the Microsoft Dynamics 365 organization. You must enter the resource URL in the following format: <code>https://<Microsoft_Dynamics_365_org_name>.api.crm8.dynamics.com</code>
OAuth Token URL	OAuth 2.0 token endpoint of the Microsoft Dynamics 365 organization. The application that is registered in Azure AD sends the access token requests to this endpoint. You must enter the following value in this field: <code>https://login.windows.net/common/oauth2/token</code>

Note: For more information about the OAuth 2.0 Username-Password Flow authentication method, see the Microsoft Dynamics 365 documentation.

Connection properties for OAuth 2.0 Client Secret Flow authentication

The following table describes the connection properties for a Microsoft Dynamics 365 Mass Ingestion connection configured with OAuth 2.0 Client Secret Flow authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Microsoft Dynamics 365 Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Client ID	Client ID of the application registered in Azure AD.
Client Secret	Client secret of the application registered in Azure AD.
Resource URL	URL of the Microsoft Dynamics 365 organization. You must enter the resource URL in the following format: <code>https://<Microsoft_Dynamics_365_org_name>.api.crm8.dynamics.com</code>
OAuth Token URL	OAuth 2.0 token endpoint of the Microsoft Dynamics 365 organization. The application that is registered in Azure AD sends the access token requests to this endpoint. You must enter the following value in this field: <code>https://login.microsoftonline.com/<tenant_id>/oauth2/token</code>

Note: For more information about the OAuth 2.0 Client Secret Flow authentication method, see the Microsoft Dynamics 365 documentation.

Connection properties for OAuth 2.0 JWT Bearer Flow authentication

The following table describes the connection properties for a Microsoft Dynamics 365 Mass Ingestion connection configured with OAuth 2.0 JWT Bearer Flow authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.

Connection property	Description
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Microsoft Dynamics 365 Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Client ID	Client ID of the application registered in Azure AD.
Certificate Signature	Base64URL string that encodes the hexadecimal value which represents the SHA-1 thumbprint of the X509 certificate.
Keystore Path	Absolute path to the keystore file that contains the X509 certificate required to validate a JSON Web Token (JWT) and establish a secure connection with Microsoft Dynamics 365. The keystore file must be in the Java KeyStore (JKS) format.
Keystore Password	Password for the keystore file.
Private Key Alias	Alias name of the private key used to sign the JWT.
Private Key Password	Password for the private key.
Audience for JWT	URL of the Microsoft Dynamics 365 resource server to which the application that is registered in Azure AD sends the JWT for validation. You must enter the address in the following format: <code>https://login.microsoftonline.com/<tenant_id>/oauth2/token</code>
Resource URL	URL of the Microsoft Dynamics 365 organization. You must enter the resource URL in the following format: <code>https://<Microsoft_Dynamics_365_org_name>.api.crm8.dynamics.com</code>
OAuth Token URL	OAuth 2.0 token endpoint of the Microsoft Dynamics 365 organization. The application that is registered in Azure AD sends the access token requests to this endpoint. You must enter the following value in this field: <code>https://login.microsoftonline.com/<tenant_id>/oauth2/token</code>

Note: For more information about the OAuth 2.0 Client Secret Flow authentication method, see the Microsoft Dynamics 365 documentation.

Microsoft SQL Server connection properties

When you set up a Microsoft SQL Server connection, configure the connection properties.

The following table describes the Microsoft SQL Server connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	Type of connection. Select SQL Server from the list.
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run a database ingestion and replication task on a Hosted Agent or serverless runtime environment.
SQL Server Version	Microsoft SQL Server database version.
Authentication Mode	<p>Authentication method to access Microsoft SQL Server.</p> <p>Select one of the following methods:</p> <ul style="list-style-type: none">- SQL Server Authentication. Uses your Microsoft SQL Server user name and password to access Microsoft SQL Server.- Windows Authentication (Deprecated). Uses the Microsoft Windows authentication to access Microsoft SQL Server. This option is available when you access Data Integration by using Microsoft Windows. When you choose this option, you don't need to enter credentials to access Microsoft SQL Server and ensure that the user account that starts the Secure Agent service is available in the Microsoft SQL Server database. Note: Windows authentication is not certified for Microsoft SQL Server 2017 version hosted on Linux. You can't configure Windows Authentication when you use a serverless runtime environment.- Active Directory Password. Uses the Azure Active Directory user name and password to authenticate and access the Microsoft Azure SQL Database. Note: Data Ingestion and Replication does not support this authentication mode.- Windows Authentication v2. Uses this authentication method to access Microsoft SQL Server from Data Integration or Data Ingestion and Replication using an agent hosted on a Linux or Windows machine. When you choose this option on Linux, enter your domain name and Microsoft Windows credentials to access Microsoft SQL Server. When you choose this option on Windows, the agent uses the user credentials specified in the connection only to test the connection. During runtime, the agent uses the credentials of the user who started the Secure Agent service. Ensure that the user account that starts the Secure Agent service is available in the Microsoft SQL Server database.- Kerberos. Uses Kerberos authentication to connect to Microsoft SQL Server. When you choose this option on Windows, ensure that the user account that starts the Secure Agent service is available in the Microsoft SQL Server database. You don't need to enter your credentials to access Microsoft SQL Server.

Property	Description
Domain	Applies to Windows Authentication v2. The domain name of the Windows user.
User Name	<p>User name for the database login. The user name can't contain a semicolon.</p> <p>To connect to Microsoft Azure SQL Database, specify the user name in the following format: username@host</p> <p>If you use Windows Authentication v2 on Windows, the user name is used as follows:</p> <ul style="list-style-type: none"> - During design time, the agent uses the user name specified here to test the connection. - During runtime, the Microsoft SQL server driver ignores the user name specified in this field and uses the credentials of the user who started the Secure Agent service. <p>If you use Windows Authentication v2 on Linux, the user name specified here is used both during design time and runtime.</p> <p>Note: This property is not applicable if you use the Windows Authentication mode to access Microsoft SQL Server.</p>
Password	<p>Password for the database login. The password can't contain a semicolon.</p> <p>If you use Windows Authentication v2 on Windows, the password is used as follows:</p> <ul style="list-style-type: none"> - During design time, the agent uses the password specified here to test the connection. - During runtime, the Microsoft SQL server driver ignores the password specified in this field and uses the credentials of the user who started the Secure Agent service. <p>If you use Windows Authentication v2 on Linux, the password specified here is used both during design time and runtime.</p> <p>Note: This property is not applicable if you use the Windows Authentication mode to access Microsoft SQL Server.</p>
Host	<p>Name of the machine hosting the database server.</p> <p>To connect to Microsoft Azure SQL Database, specify the fully qualified host name.</p> <p>For example, vmjcmwxsfbheng.westus.cloudapp.azure.com.</p>
Port	<p>Network port number used to connect to the database server.</p> <p>Default is 1433.</p>
Instance Name	Instance name of the Microsoft SQL Server database.
Database Name	<p>Database name for the Microsoft SQL Server target connection. Database name is case-sensitive if the database is case-sensitive. Maximum length is 100 characters.</p> <p>Database names can include alphanumeric and underscore characters.</p>
Schema	Schema used for the target connection.
Code Page	The code page of the database server.
Encryption Method	<p>The method that the Secure Agent uses to encrypt the data sent between the driver and the database server. You can use the encryption method to connect to Microsoft Azure SQL Database.</p> <p>Default is None.</p>
Crypto Protocol Version	Cryptographic protocols to use when you enable SSL encryption.

Property	Description
Validate Server Certificate	<p>When set to True, Secure Agent validates the certificate that is sent by the database server. If you specify the HostNameInCertificate parameter, Secure Agent also validates the host name in the certificate.</p> <p>When set to false, the Secure Agent doesn't validate the certificate that is sent by the database server.</p>
Trust Store	The location and name of the truststore file. The truststore file contains a list of Certificate Authorities (CAs) that the driver uses for SSL server authentication.
Trust Store Password	The password to access the contents of the truststore file.
Host Name in Certificate	Host name of the machine that hosts the secure database. If you specify a host name, the Secure Agent validates the host name included in the connection with the host name in the SSL certificate.
Metadata Advanced Connection Properties	<p>Additional properties for the JDBC driver to fetch the metadata.</p> <p>If you specify more than one property, separate each key-value pair with a semicolon.</p> <p>For example,</p> <pre>LoginTimeout=100</pre> <p>Note: The default connection timeout is 270 seconds. Add this property to configure the connection timeout during design time.</p>
Runtime Advanced Connection Properties	<p>Additional properties for the ODBC driver to run mappings.</p> <p>If you specify more than one property, separate each key-value pair with a semicolon.</p> <p>For example,</p> <pre>LoginTimeout=100</pre> <p>Note: The default connection timeout is 270 seconds. Add this property to configure the connection timeout during runtime.</p>

Microsoft Fabric OneLake connection properties

When you set up a Microsoft Fabric OneLake connection, configure the connection properties.

The following table describes the Microsoft Fabric OneLake connection properties:

Property	Description
Connection Name	<p>Name of the connection.</p> <p>Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.</p>
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Microsoft Fabric OneLake connection type.

Property	Description
Runtime Environment	The name of the runtime environment where you want to run the tasks. Select a Secure Agent, Hosted Agent, or serverless runtime environment. Note: The Hosted Agent and serverless runtime environments are not supported by application ingestion and replication, database ingestion and replication, and file ingestion and replication tasks.
Workspace Name	Name of the workspace in Microsoft Fabric OneLake.
Lakehouse Path	Path or name of the lakehouse present in the workspace. You can specify the path in one of the following ways: - <i>root directory (/)</i> to access the files in the workspace. - <i>lakehouse name/Files</i> to access the files present in the lakehouse.
Authentication Type	Authentication type to access Microsoft Fabric OneLake. Service Principal Authentication uses the client ID, client secret, and tenant ID to connect to Microsoft Fabric OneLake.
Client ID	The application ID or client ID of your application registered in the Azure Active Directory.
Client Secret	The client secret of your application registered in the Azure Active Directory.
Tenant ID	The ID of the Azure Active Directory instance in which you created the application.
Microsoft Fabric OneLake Endpoint	The type of Microsoft Fabric OneLake endpoint that you want to connect to. Default is fabric.microsoft.com .

MongoDB Mass Ingestion connection properties

When you set up a MongoDB Mass Ingestion connection, you must configure the connection properties.

The following table describes MongoDB Mass Ingestion connection properties:

Connection property	Description
Connection Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 255 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
Description	Optional. A description of the connection. The description cannot exceed 4,000 characters.
Type	Type of connection. You must select MongoDB Mass Ingestion .
Runtime Environment	The name of the runtime environment where you want to run the tasks.

Connection property	Description
Host and Port	An SRV record or a comma-separated list of <i>host_name:port</i> pairs. Note: If you are using the MongoDB replica set mode, you can enter multiple host names for resilience. If one host is not available, another specified host will be used.
SRV	Select this check box if you specified an SRV record in Host and Port property.
User Name	User name for logging in to the database.
Password	Password for the specified database user.
Authentication Database	The name of the authentication database associated with the specified user.
Replica Set Name	The name of the replica set that is composed of the MongoDB servers with replicas of the source data. This field is relevant if you are using the MongoDB replica set mode.
Additional Connection Properties	One or more additional MongoDB connection string options that you want to use. Specify the properties as key-value pairs. If you specify more than one property, separate them with the ampersand symbol (&). The connection properties are case sensitive. Example: <code>authSource=admin&replicaSet=rsprimary</code> For more information about the MongoDB connection string options, refer to: https://www.mongodb.com/docs/v5.2/reference/connection-string/#connection-string-options

MQTT connection properties

When you set up an MQ Telemetry Transport (MQTT) connection, you must configure the connection properties.

The following table describes the MQTT connection properties:

Property	Description
Connection Name	Name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: <code>~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /</code>
Description	Optional. Description that you can use to identity the connection. The description cannot exceed 4,000 characters.
Type	The MQTT connection type. If you do not see the connection type, go to the Add-On Connectors page to install the connector.

Property	Description
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Broker URI	The connection URL of the MQTT broker. If specified, this value overrides the URL specified in the main portion of the URL. Sample URL: <code>tcp://<IP Address>:<port></code>
Client Id	Client identifier of your MQTT client. If this value is left blank, the MQTT server assigns a unique value. This property value must be unique for each MQTT client connecting to a specific MQTT server. Sharing projects without changing the Client ID can lead to connection issues, including disconnections and missing updates.
Username	Username to use when connecting to the broker.
Password	Password to use when connecting to the broker.
Connection Timeout	Maximum time interval the client will wait for the connection to the MQTT server to be established. Default timeout is 30 seconds. A value of 0 disables timeout processing. That is, the client waits until the network connection is made successfully or fails.
Use SSL	Enable this option to use SSL for secure transmission. If you enable the SSL authentication, ensure to provide both keystore and truststore details for using the MQTT connection in a streaming ingestion and replication task.
Keystore Filename	Contains the keys and certificates required for secure communication.
Keystore Password	Password for the keystore filename.
Keystore Type	Type of keystore to use. Keystore type defines the storage and data format of the keystore information and the algorithms used to protect private keys in the keystore. Use one of the following types: <ul style="list-style-type: none"> - JKS. Stores private keys and certificates. - PKCS12. Stores private keys, secret keys. and certificates.
Truststore Filename	File name of the truststore file.
Truststore Password	Password for the truststore file name.

Property	Description
Truststore Type	Type of truststore to use. Use one of the following types: <ul style="list-style-type: none"> - JKS - PKCS 12
TLS Protocol	Transport protocols to use. Use one of the following types: <ul style="list-style-type: none"> - SSL - SSLv3 - TLS - TLSv1 - TLSv1.1 - TLSv1.2

MySQL connection properties

When you set up a MySQL connection, configure the connection properties.

The following table describes the MySQL connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	Type of connection. Select MySQL from the list.
Runtime Environment	The name of the runtime environment where you want to run the tasks. Note: You cannot run a database ingestion and replication task on a Hosted Agent or serverless runtime environment.
User Name	User name for the database login. The user name can't contain a semicolon.
Password	Password for the database login. The password can't contain a semicolon.
Host	Name of the machine that hosts the database server.
Port	Network port number used to connect to the database server. Default is 3306.
Database Name	Name of the MySQL database that you want to connect to. Note: The database name is case-sensitive. Maximum length is 64 characters. Database name can contain alphanumeric and underscore characters.

Property	Description
Code Page	The code page of the database server.
Metadata Advanced Connection Properties	<p>Additional properties for the JDBC driver to fetch the metadata. Enter properties in the following format:</p> <pre><parameter name>=<parameter value></pre> <p>If you enter more than one property, separate each key-value pair with a semicolon.</p> <p>For example, enter the following property to configure the connection timeout when you test a connection:</p> <pre>connectTimeout=<value_in_milliseconds></pre> <p>Note: The default connection timeout is 270000 milliseconds.</p>
Runtime Advanced Connection Properties	<p>Additional properties for the ODBC driver to run ingestion and replication jobs.</p> <p>If you specify more than one property, separate each key-value pair with a semicolon.</p>

Netezza connection properties

When you set up a Netezza connection, you must configure the connection properties.

The following table describes the Netezza connection properties:

Connection property	Description
Connection Name	<p>Name of the connection.</p> <p>Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -,</p> <p>Maximum length is 255 characters.</p>
Description	Description of the connection. Maximum length is 4000 characters.
Type	Netezza.
Runtime Environment	<p>The name of the runtime environment where you want to run tasks.</p> <p>Specify a Secure Agent.</p>
Database	The name of the Netezza database.
Schemaname	<p>The schema used for the Netezza source or target.</p> <p>Schema name is case sensitive.</p>
Servename	The Netezza database host name.
Port	<p>Network port number used to connect to the database server.</p> <p>Default is 1521.</p>
Driver	The Netezza ODBC driver name, NetezzaSQL, used to connect to the Netezza database.

Connection property	Description
Runtime Additional Connection Configuration	Additional run-time attributes required to fetch data. For example, <code>securityLevel=preferredUnSecured;caCertFile =</code>
Metadata Additional Connection Configuration	The values to set the optional properties of the JDBC driver to fetch the metadata.
Username	Database user name with the appropriate read and write database permissions to access the database.
Password	Password for the database user name.

NetSuite Mass Ingestion connection properties

When you set up a NetSuite Mass Ingestion connection, you must configure the connection properties.

Note: Before you configure the connection properties, install the SuiteAnalytics Connect JDBC driver and copy the NQjc.jar file to the following directory: `<Secure_Agent>\ext\connectors\thirdparty\informatica.netsuiteami`

For more information about installing the SuiteAnalytics Connect JDBC driver, see the [SuiteAnalytics Connect documentation](#).

The following table describes the connection properties for a NetSuite Mass Ingestion connection:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: <code>_ . + -</code> . Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Select the Netsuite Mass Ingestion connection type.
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Email ID	User name of the NetSuite account. The user name is an email address.
Password	Password for the NetSuite account.

Connection property	Description
Service Host	<p>Name of the SuiteAnalytics Connect Service host.</p> <p>The value in this field must match the value specified in the Service Host field under the Your Configuration section of the SuiteAnalytics Connect Driver Download page in NetSuite. To access the SuiteAnalytics Connect Driver Download page, log in to NetSuite and click the Set Up SuiteAnalytics Connect link in the Settings portlet.</p>
Service Port	TCP/IP port on which the SuiteAnalytics Connect server is listening. Default is 1708.
Service Datasource	<p>Data source that you want to use to access NetSuite data. You can select one of the following data sources:</p> <ul style="list-style-type: none"> - NetSuite.com - NetSuite2.com <p>Default is NetSuite2.com.</p> <p>Note:</p> <ul style="list-style-type: none"> - In connections configured before the August 2022 release, the default value for this field is NetSuite.com. - To use a NetSuite2.com data source, the NetSuite user account must be configured with some specific roles and permissions. For more information about the roles and permissions required to access NetSuite2.com data sources, see the NetSuite documentation.
Account ID	<p>NetSuite account ID.</p> <p>To find your account ID, log in to NetSuite and click Setup > Integration > Web Services Preferences.</p> <p>If you cannot access the Setup menu, navigate to Support > Go to Suite Answers > Contact support by phone. The page displays your account ID.</p>
Role ID	Role ID associated with the NetSuite account.
Additional Connection Properties	<p>Additional properties for the SuiteAnalytics Connect Driver that is used to connect to the NetSuite service data source. Specify the properties in <property>=<value> format. If you want to specify multiple properties, separate each property-value pair with a semicolon (;).</p> <p>You can specify the following connection properties in this field:</p> <ul style="list-style-type: none"> - ValidateServerCertificate: Determines whether the driver validates the certificate sent by the SuiteAnalytics Connect server. During SSL server authentication, the SuiteAnalytics Connect server sends a certificate issued by a trusted Certificate Authority (CA). The required CAs are usually included in the Java truststore but you can also specify them using the TrustStore property. Valid values for the ValidateServerCertificate property are <i>true</i> and <i>false</i>. - TrustStore: Contains the path to a valid truststore containing the security certificates to be used for server authentication. The TrustStore property is ignored if the ValidateServerCertificate property is set to <i>false</i>. <p>Note: For more information about the additional connection properties, see the NetSuite documentation.</p>

OPC UA connection properties

When you set up an OPC UA connection, you must configure the connection properties.

The following table describes the OPC UA connection properties:

Property	Description
Connection Name	Name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
Description	Optional. Description of the connection. The description cannot exceed 4,000 characters.
Type	The OPC UA connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks.
Endpoint URL	The unique URL to connect to the OPC UA server. The endpoint URL identifies the specific instance of a server and a security policy type. A valid endpoint URL consists of the endpoint type (opc.tcp), the endpoint host name (IP address, URL, or DSN), and the endpoint port number. For example, <code>opc.tcp://opcuaserver.com:48010</code>
Security Policy	The security policy used to connect to the OPC UA server. The security policy parameters specify the security algorithms that the OPC UA server supports. You can choose one of the following security policies: <ul style="list-style-type: none">- None. No security provided.- Basic128Rsa15- Basic256- Basic256Sha256- Aes128_Sha256_RsaOaep- Aes256_Sha256_RsaPss Note: The OPC Foundation deprecated the security policies, Basic128Rsa15 and Basic256 as of OPC UA specification version 1.04. The encryption provided by these policies is less secure. Use these security policies only to provide backward compatibility.
Security Mode	The security mode used to connect to the OPC UA server. The security mode is valid only when security policy is not set to None. You can choose one of the following security policies: <ul style="list-style-type: none">- Sign. Transfer unencrypted data, but with digital signatures that allow verification of data integrity.- SignAndEncrypt. Transfer signed and encrypted data.

Property	Description
Application URI	<p>Optional. A unique identifier that the OPC UA application can use to connect to the OPC UA server.</p> <p>Enter a unique ID in the following format:</p> <pre>urn:aaa:bbb</pre> <p>For example, <code>urn:nifi:opcua</code></p> <p>The unique identifier must match the URI of the Subject Alternative Name of your OPC UA client certificate.</p>
Client Keystore Location	<p>Optional. Absolute path and file name of the keystore file that contains private keys and certificates for the OPC UA server.</p> <p>Enter the path in the following format:</p> <pre>/root/opcua/client.jks</pre> <p>The keystore must contain only one keypair entry of private key and certificate. If multiple keypair entries exist, the first entry is used.</p>
Client Keystore Password	Optional. Password for the client keystore.
Require server authentication	Optional. Enable if you require server authentication of client certificates, client authentication of server certificates, or both.
Trust store Location	<p>Optional. The absolute path of the truststore file that contains the trusted certificate.</p> <p>Enter the path in the following format:</p> <pre>/root/opcua/trust.jks</pre>
Trust store Password	Password for the truststore file.
Authentication Policy	<p>Authentication settings required to establish the connections.</p> <p>You can choose one of the following authentication policies:</p> <ul style="list-style-type: none"> - Anon. Anonymous authentication. Anonymous tokens are associated with servers that do not require user authentication. - UserName. User name and password tokens are associated with servers with any password based system, such as Windows.
User Name	User name to access the OPC UA server if you choose authentication policy as UserName .
Password	Password to access the OPC UA server if you choose authentication policy as UserName .

Oracle Cloud Object Storage connection properties

When you define an Oracle Cloud Object Storage connection for a database ingestion and replication task, you must configure connection properties.

The following table describes the Oracle Cloud Object Storage connection properties:

Property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: <code>_ . + - ,</code> . Maximum length is 255 characters.
Description	An optional description for the connection. Maximum length is 4000 characters.
Type	The type of connection. For an Oracle Cloud Object Storage Database Ingestion connection, the type must be Oracle Cloud Object Storage .
Runtime Environment	The name of the runtime environment where you want to run database ingestion and replication tasks. You define runtime environments in Administrator.
Authentication Type	The authentication type that is used to connect to Oracle Cloud Object Storage to stage the files. Select one of the following options: <ul style="list-style-type: none">- Simple Authentication. API key-based authentication.- ConfigFile Authentication. Identity credential details are provided through a configuration file.
User	If you selected Simple Authentication as the authentication type, specify the Oracle Cloud Identifier (OCID) of the user for whom the key pair is added.
Fingerprint	If you selected Simple Authentication as the authentication type, specify the fingerprint of the public key.
Tenancy	If you selected Simple Authentication as the authentication type, specify the Oracle Cloud Identifier (OCID) of the tenancy that is the globally unique name of the OCI account.
PrivateKey File Location	If you selected Simple Authentication as the authentication type, specify the location of the private key file in .PEM format on the Secure Agent machine.
Config File Location	If you selected ConfigFile Authentication as the authentication type, specify the location of the configuration file on the Secure Agent machine. Enter the absolute path. If you do not enter any value, <code><System default location>/oci/config</code> is used to retrieve the configuration file.
Profile Name	If you selected ConfigFile Authentication as the authentication type, specify the name of the profile in the configuration file that you want to use. Default is <code>DEFAULT</code> .
Bucket Name	The Oracle Cloud Storage bucket name. The bucket contains the objects and files.

Property	Description
Folder Path	The path to the folder under the specified Oracle Cloud Storage bucket. For example, bucket/Dir_1/Dir_2/FileName.txt. Here, Dir_1/Dir_2 is the folder path.
Region	The Oracle Cloud Object Storage region where the bucket exists.

Oracle Database Ingestion connection properties

When you define an Oracle Database Ingestion connection for a database ingestion and replication task, you must configure connection properties.

The following table describes the connection properties:

Property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Oracle Database Ingestion .
Runtime Environment	The name of the runtime environment where you want to run database ingestion and replication tasks. You define runtime environments in Administrator.
User Name	User name for the Oracle database login. The user name cannot contain a semicolon.
Password	Password for the Oracle database login. The password cannot contain a semicolon.
Host	Host name of the database server.
Port	Network port number used to connect to the database server. Default is 1521.
Service Name	Service name or System ID (SID) that uniquely identifies the Oracle database. Specify the SID in the following format, including the leading semicolon (;), to connect to Oracle databases: ;SID=<ORACLE_SID>
Schema	Schema used for the Oracle connection.

Property	Description
Code Page	The code page of the database server. Database ingestion and replication tasks use the UTF-8 code page. Default is UTF-8.
Encryption Method	<p>For initial load jobs, determines whether the data exchanged between the Secure Agent and the Oracle database server is encrypted:</p> <p>Options are:</p> <ul style="list-style-type: none"> - SSL. Establishes a secure connection using SSL for data encryption. If the Oracle database server cannot configure SSL, the connection fails. - No Encryption. Establishes a connection without using SSL. Data is not encrypted. <p>Default is No Encryption.</p>
Crypto Protocol Version	<p>If you selected SSL as the encryption method, you must specify a cryptographic protocol or a list of cryptographic protocols supported by your server to use with an encrypted connection. Options are:</p> <ul style="list-style-type: none"> - SSLv2 - SSLv3 - TLSv1.2 <p>Default is TLSv1.2.</p>
Validate Server Certificate	<p>If you selected SSL as the encryption method, controls whether the Secure Agent validates the server certificate that is sent by the Oracle database server.</p> <ul style="list-style-type: none"> - True. Validate the server certificate. - False. Do not validate the server certificate. <p>Default is False.</p> <p>If you also specify the Host Name in Certificate property, the Secure Agent also validates the host name in the certificate.</p>
Trust Store	If you selected SSL as the encryption method and enabled validation of the server certificate, specify the path and name of the truststore file, which contains the list of the Certificate Authorities (CAs) that the client trusts for SSL authentication.
Trust Store Password	If you selected SSL as the encryption method and enabled validation of the server certificate, specify a password for accessing the contents of the truststore file.
Host Name in Certificate	If you selected SSL as the encryption method and enabled validation of the server certificate, specify the host name of the machine that hosts the Oracle database to provide for additional security. The Secure Agent validates the host name included the connection with the host name in the SSL certificate.

Property	Description
Key Store	If you selected SSL as the encryption method and client authentication is enabled on the Oracle database server, specify the path and name of the keystore file. The keystore file contains the certificates that the client sends to the Oracle server in response to the server's certificate request.
Key Store Password	If you selected SSL as the encryption method and client authentication is enabled on the Oracle database server, specify the password for the keystore file.
Key Password	If you selected SSL as the encryption method and client authentication is enabled on the Oracle database server, specify the password for the keys in the keystore file. Use this property when the keys have a different password than the keystore file.
Database Connect String	A TNS name, an Oracle Net keyword-value pair, or a SQL connect string URL that OCI uses to connect to Oracle.
TDE Wallet Directory	The path to the directory that contains the Oracle wallet file used for Oracle Transparent Data Encryption (TDE). Specify this property value only if you capture change data from TDE-encrypted tablespaces and one of the following conditions are true: <ul style="list-style-type: none"> - The Oracle wallet is not available to the database. - The Oracle database is running on a server that is remote from Oracle redo logs. - The wallet directory is not in the default location on the database host or the wallet name is not the default name of ewallet.p12. - The wallet directory is not available to the Secure Agent host.
TDE Wallet Password	A clear text password that is required to access the Oracle TDE wallet and get the master key. This property value is required if you need to read and decrypt data from TDE-encrypted tablespaces in the Oracle source database.

Property	Description
Directory Substitution	<p>A local path prefix to substitute for the server path prefix of the redo logs on the Oracle server. This substitute local path is required when the log reader runs on a system other than the Oracle server and uses a different mapping to access the redo log files. Use this property in the following situations:</p> <ul style="list-style-type: none"> - The redo logs reside on shared disk. - The redo logs have been copied to a system other than the Oracle system. - The archived redo logs are accessed by using a different NFS mount. <p>Do <i>not</i> use this statement if you use Oracle Automatic Storage Management (ASM) to manage the redo logs. You can define one or more substitutions in the following format:</p> <pre>server_path_prefix,local_path_prefix;server_path_prefix,local_path_prefix;...</pre> <p>Note: This property does not apply to Oracle targets.</p>
Reader Active Log Mask	<p>A mask that the log reader uses for selecting active redo logs when the Oracle database uses multiplexing of redo logs. The log reader compares the mask against the member names in an active redo log group to determine which log to read. In the mask, you can use the asterisk (*) wildcard to represent zero or more characters.</p> <p>The mask can be up to 128 characters in length. It is case-sensitive on Linux or UNIX systems but not on Windows systems.</p> <p>Note: This property does not apply to Oracle targets.</p>
Reader Archive Destination 1	<p>The primary log destination from which the log reader reads archived logs, when Oracle is configured to write more than one copy of each archived redo log. Enter a number that corresponds to a <i>n</i> value in an Oracle LOG_ARCHIVE_DEST_<i>n</i> initialization parameter, where <i>n</i> is a value from 1 to 10.</p> <p>If you set only one of the Reader Archive Destination 1 and Destination 2 properties, the log reader uses that property setting. If you specify neither property, the archive log queries are not filtered by the log destination.</p> <p>Note: This property does not apply to Oracle targets.</p>
Reader Archive Destination 2	<p>The secondary log destination from which the log reader reads archived logs when the primary destination becomes unavailable or when the logs at the primary destination cannot be read. For example, logs might have been corrupted or deleted. Enter a number that corresponds to the <i>n</i> value in an Oracle LOG_ARCHIVE_DEST_<i>n</i> initialization parameter, where <i>n</i> is a value from 1 to 10. Usually, this value is a number greater than 1.</p> <p>Note: This property does not apply to Oracle targets.</p>

Property	Description
Reader ASM Connect String	In an Oracle ASM environment, the Oracle connection string, defined in TNS, that the log reader uses to connect to the ASM instance that manages storage of active and archived redo logs for the source database. Note: This property does not apply to Oracle targets.
Reader ASM User Name	In an Oracle ASM environment, an Oracle user ID that the log reader uses to connect to the ASM instance that manages storage of active and archived redo logs for the source database. This user ID must have SYSDBA or SYSASM authority. To use SYSASM authority, set the Reader ASM Connect As SYSASM property to Y. Note: This property does not apply to Oracle targets.
Reader ASM Password	In an Oracle ASM environment, a clear text password for the user that is specified in the Reader ASM User Name property. The log reader uses this password and the ASM user name to connect to the ASM instance that manages storage of active and archived redo logs for the source database. Note: This property does not apply to Oracle targets.
Reader ASM Connect As SYSASM	If you use Oracle 11g ASM or later and want the log reader to use a user ID that has SYSASM authority to connect to the ASM instance, select this check box. Also specify a user ID that has SYSASM authority in the Reader ASM User Name property. To use a user ID that has SYSDBA authority, clear this check box. By default, this check box is cleared. Note: This property does not apply to Oracle targets.

Property	Description
Reader Mode	<p>Indicates the source of and types of Oracle redo logs that the log reader reads. Valid options are:</p> <ul style="list-style-type: none"> - ACTIVE. Read active and archived redo logs from the Oracle online system. Optionally, you can use the Reader Active Log Mask property to filter the active redo logs and use the Reader Archive Destination 1 and Reader Archive Destination 2 properties to limit the archived log destinations from which to read archived logs. - ARCHIVEONLY. Read only archived redo logs. Optionally, you can use the Reader Archive Destination 1 and Reader Archive Destination 2 properties to limit the archived log destinations from which to read archived logs. - ARCHIVECOPY. Read archived redo logs that have been copied to an alternate file system. For combined initial and incremental load jobs, you must also set the source custom property <code>pxw.cdcreader.oracle.reader.additional</code> with the dir and file parameters, at the direction of Informatica Global Customer Support. <p>You can use this option in the following situations:</p> <ul style="list-style-type: none"> - You do not have the authority to access the Oracle archived redo logs directly. - The archived redo logs are written to ASM, but you do not have access to ASM. - The archived log retention policy for the database server causes the archived logs to not be retained long enough. <p>With this option, the Reader Archive Destination 1 and Reader Archive Destination 2 properties are ignored.</p> <p>Default is ACTIVE.</p> <p>Note: This property does not apply to Oracle targets.</p>
Reader Standby Log Mask	<p>A mask that the log reader uses for selecting redo logs for an Oracle physical standby database when the database uses multiplexing of redo logs. The log reader compares the mask against the member names in an redo log group to determine which log to read. In the mask, you can use the asterisk (*) wildcard to represent zero or more characters.</p> <p>The mask can be up to 128 characters in length. It is case-sensitive on Linux or UNIX systems but not on Windows systems.</p> <p>Note: This property does not apply to Oracle targets.</p>
Standby Connect String	<p>An Oracle connection string, defined in TNS, that the log reader uses to connect to the Oracle physical standby database for change capture when the database is not open with read only access.</p> <p>Note: This property does not apply to Oracle targets.</p>
Standby User Name	<p>A user ID that the log reader uses to connect to the Oracle physical standby database for change capture. This user ID must have SYSDBA authority.</p> <p>Note: This property does not apply to Oracle targets.</p>

Property	Description
Standby Password	A password that the log reader uses to connect to the Oracle physical standby database for change capture. Note: This property does not apply to Oracle targets.
RAC Members	The maximum number of active redo log threads, or <i>members</i> , in an Oracle Real Application Cluster (RAC) that can be tracked. For a Data Guard physical standby database that supports a primary database in a RAC environment, this value is the number of active threads for the primary database. Valid values are 1 to 100. Default is 0, which causes an appropriate number of log threads to be determined automatically. If this value is not appropriate for your environment, set this property to a value greater than 0. Note: This property does not apply to Oracle targets.
BFILE Access	Select this check box in the following circumstances: <ul style="list-style-type: none"> - You use BFILE access to redo logs in physical directories on the local Oracle server file system. BFILE access uses Oracle directory objects to remotely access the redo logs in the file system. This method is an alternative to other log access methods such as ASM or NFS mounts. - You have an Amazon Relational Database Service (RDS) for Oracle source. In this case, this option enables access to the redo logs of a cloud-based database instance deployed in RDS. By default, this check box is cleared. Note: This property does not apply to Oracle targets.

Oracle Fusion Cloud Mass Ingestion connection properties

When you set up an Oracle Fusion Cloud Mass Ingestion connection, you must configure the connection properties.

Note: Oracle Fusion Cloud Mass Ingestion connections can access the data of only Enterprise Resource Planning (ERP), Human Capital Management (HCM), and Oracle Supply Chain and Manufacturing (SCM) modules of Oracle Fusion Cloud Applications Suite.

The following table describes the connection properties for an Oracle Fusion Cloud Mass Ingestion connection:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. Select the Oracle Fusion Cloud Mass Ingestion connection type.
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Authentication	Authentication method of the connection. By default, the connection uses the Basic authentication method.
User Name	User name of the Oracle Cloud account.
Password	Password for the Oracle Cloud account.
Server URL	URL of the Oracle Cloud service that you want to access.
API Version	Version of the Oracle Cloud REST API that you want to use for the connection. Optional for the BICC replication approach.

PostgreSQL connection properties

When you set up a PostgreSQL connection, configure the connection properties.

The following table describes the PostgreSQL connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	Type of connection. Select PostgreSQL from the list.

Property	Description
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run a database ingestion and replication task on a Hosted Agent or in a serverless runtime environment.
Host Name	Host name of the PostgreSQL server to which you want to connect.
Port	Port number for the PostgreSQL server to which you want to connect. Default is 5432.
Schema	The schema name. If you don't specify the schema name, all the schemas available in the database are listed while importing the source object in Data Integration.
Database	The PostgreSQL database name.
User Name	User name to access the PostgreSQL database.
Password	Password for the PostgreSQL database user name.
Encryption Method	Determines whether the data exchanged between the Secure Agent and the PostgreSQL database server is encrypted. Select one of the following encryption methods: <ul style="list-style-type: none"> - noEncryption. Establishes a connection without using SSL. Data is not encrypted. - SSL. Establishes a connection using SSL. Data is encrypted using SSL. If the PostgreSQL database server can't configure SSL, the connection fails. - requestSSL. Attempts to establish a connection using SSL. If the PostgreSQL database server can't configure SSL, the Secure Agent establishes an unencrypted connection. Default is noEncryption.
Validate Server Certificate	Applicable if you select SSL or requestSSL as the encryption method. Select the Validate Server Certificate option so that the Secure Agent validates the server certificate that is sent by the PostgreSQL database server. If you specify the Host Name In Certificate property, the Secure Agent also validates the host name in the certificate.
TrustStore	Applicable if you select SSL or requestSSL as the encryption method and the Validate Server Certificate option. The path and name of the truststore file, which contains the list of the Certificate Authorities (CAs) that the PostgreSQL client trusts.
TrustStore Password	Applicable if you select SSL or requestSSL as the encryption method and the Validate Server Certificate option. The password to access the truststore file that contains the SSL certificate.
Host Name In Certificate	Optional when you select SSL or requestSSL as the encryption method and the Validate Server Certificate option. A host name for providing additional security. The Secure Agent validates the host name included in the connection with the host name in the SSL certificate.
KeyStore	Applicable if you select SSL as the encryption method and when client authentication is enabled on the PostgreSQL database server. The path and the file name of the key store. The keystore file contains the certificates that the PostgreSQL client sends to the PostgreSQL server in response to the server's certificate request.

Property	Description
KeyStore Password	Applicable if you select SSL as the encryption method and when client authentication is enabled on the PostgreSQL database server. The password for the keystore file required for secure communication.
Key Password	Applicable if you select SSL as the encryption method and when client authentication is enabled on the PostgreSQL database server. Required when individual keys in the keystore file have a different password than the keystore file.
Additional Connection Properties	Additional connection parameters that you want to use. Provide the connection parameters as semicolon-separated key-value pairs.
Crypto Protocol Versions	Required if you select SSL or requestSSL as the encryption method. A cryptographic protocol or a list of cryptographic protocols to use with an encrypted connection. You can select one of the following protocols: <ul style="list-style-type: none"> - SSLv3 - TLSv1_2 Default is TLSv1_2.

REST V2 connection properties

When you set up a REST V2 connection, you must configure the connection properties.

The following table describes the REST V2 connection properties for a Standard authentication type connection:

Connection property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The REST V2 connection type.
Runtime Environment	Name of the runtime environment where you want to run the tasks. Select a Secure Agent, Hosted Agent, or serverless runtime environment. You cannot run a streaming ingestion and replication task on a Hosted Agent or serverless runtime environment.
Authentication	The authentication method that the REST V2 Connector must use to connect to the REST endpoint. Select Standard .

Connection property	Description
Authentication Type	<p>The authentication type that you can use when you select the Standard authentication.</p> <p>You can select one of the following authentication types:</p> <ul style="list-style-type: none"> - BASIC - DIGEST - OAUTH - NONE <p>Default is NONE.</p>
Auth User ID	<p>The user name to log in to the web service application when you select the Standard authentication.</p> <p>Digest authentication is not applicable.</p>
Auth Password	<p>The password associated with the user name when you select the Standard authentication.</p> <p>Digest authentication is not applicable.</p>
OAuth Consumer Key	<p>The client key associated with the web service application.</p> <p>Required only for OAuth authentication type.</p>
OAuth Consumer Secret	<p>The client password to connect to the web service application.</p> <p>Required only for OAuth authentication type.</p>
OAuth Token	<p>The access token to connect to the web service application.</p> <p>Required only for OAuth authentication type.</p>
OAuth Token Secret	<p>The password associated with the OAuth token.</p> <p>Required only for OAuth authentication type.</p>
Swagger File Path	<p>The path of the Swagger file or OpenAPI file.</p> <p>You can specify one of the following file paths:</p> <ul style="list-style-type: none"> - Absolute path along with the file name - Hosted URL <p>If you provide the absolute path of the Swagger file or OpenAPI file, the file must be located on the Secure Agent machine.</p> <p>The hosted URL must return the content of the file without prompting for further authentication and redirection.</p> <p>For example, the path of the Swagger file can be:</p> <p>C:\Swagger\sampleSwagger.json</p> <p>The user must have the read permission for the folder and the file.</p>
TrustStore File Path	<p>The absolute path of the truststore file that contains the TLS certificate to establish a one-way or two-way secure connection with the REST API. Specify a directory path that is available on each Secure Agent machine.</p> <p>You can also configure the truststore file name and password as a JVM option or import the certificate to the following directory:</p> <p><Secure Agent installation directory>\jre\lib\security\cacerts.</p> <p>For the serverless runtime environment, specify the truststore file path in the serverless agent directory.</p> <p>For example, /home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</p>

Connection property	Description
TrustStore Password	The password for the truststore file that contains the SSL certificate. You can also configure the truststore password as a JVM option.
KeyStore File Path	The absolute path of the keystore file that contains the keys and certificates required to establish a two-way secure communication with the REST API. Specify a directory path that is available on each Secure Agent machine. You can also configure the keystore file name and location as a JVM option or import the certificate to any directory. For the serverless runtime environment, specify the keystore file path in the serverless agent directory. For example, <code>/home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</code>
KeyStore Password	The password for the keystore file required for secure communication. You can also configure the keystore password as a JVM option.
Proxy Type	Type of proxy. You can select one of the following options: <ul style="list-style-type: none"> - No Proxy. Bypasses the proxy server configured in the agent or the connection properties. - Platform Proxy. Considers the proxy configured in the agent. - Custom Proxy. Considers the proxy configured in the connection properties.
Proxy Configuration	The format required to configure proxy. You can configure proxy using the following format: <code><host>:<port></code> You cannot configure an authenticated proxy server.
Advanced Fields	Enter the arguments that the agent uses when connecting to a REST endpoint. You can specify the following arguments, each separated by a semicolon (;): <ul style="list-style-type: none"> - ConnectionTimeout. The wait time in milliseconds to get a response from a REST endpoint. The connection ends after the connection timeout is over. Default is the timeout defined in the endpoint API. - Note: If you define both the REST V2 connection timeout and the endpoint API timeout, the connection ends at the shortest defined timeout. - connectiondelaytime. The delay time in milliseconds to send a request to a REST endpoint. Default is 10000. - retryattempts. Number of times the connection is attempted when 400 and 500 series error codes are returned in the response. Default is 3. Specify 0 to disable the retry attempts. - qualifiedSchema. Determines if the schema selected is qualified or unqualified. Default is false. For example, <code>connectiondelaytime:10000;retryattempts:5</code> Note : In a streaming ingestion and replication task, only <code>ConnectionTimeout</code> and <code>retryattempts</code> are applicable.

Table 1. For OAuth 2.0-Client Credentials authentication

Connection property	Description
Access Token URL	Access token URL configured in your application.
Client ID	Client ID of your application.

Connection property	Description
Client Secret	Client secret of your application.
Scope	Specifies access control if the API endpoint has defined custom scopes. Enter space separated scope attributes. For example: root_readonly root_readwrite manage_app_users
Access Token Parameters	Additional parameters to use with the access token URL. Parameters must be defined in the JSON format. For example, [{"Name": "resource", "Value": "https://<serverName>"}]
Client Authentication	Select an option to send Client ID and Client Secret for authorization either in the request body or in the request header. Default is Send Client Credentials in Body .
Generate Access Token	Generates access token based on the information provided in the above fields.
Access Token	Enter the access token value or click Generate Access Token to populate the access token value. To pass the generate access token call through a proxy server, you must configure an unauthenticated proxy server at the Secure Agent level. The REST V2 connection-level proxy configuration does not apply to the generate access token call.
Swagger File Path	The path of the Swagger file or OpenAPI file. You can specify one of the following file paths: - Absolute path along with the file name - Hosted URL If you provide the absolute path of the swagger file or OpenAPI file, the file must be located on the Secure Agent machine. The hosted URL must return the content of the file without prompting for further authentication and redirection. For example, the path of the swagger file can be: C:\swagger\sampleSwagger.json The user must have the read permission for the folder and the file. Note: In a streaming ingestion and replication task, use only a hosted URL of the swagger specification file as the swagger file path.
TrustStore File Path	The absolute path of the truststore file that contains the TLS certificate to establish a one-way or two-way secure connection with the REST API. Specify a directory path that is available on each Secure Agent machine. You can also configure the truststore file name and password as a JVM option or import the certificate to the following directory: <Secure Agent installation directory>\jre\lib\security\cacerts. For the serverless runtime environment, specify the truststore file path in the serverless agent directory. For example, /home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks
TrustStore Password	The password for the truststore file that contains the SSL certificate. You can also configure the truststore password as a JVM option.

Connection property	Description
KeyStore File Path	<p>The absolute path of the keystore file that contains the keys and certificates required to establish a two-way secure communication with the REST API. Specify a directory path that is available on each Secure Agent machine.</p> <p>You can also configure the keystore file name and location as a JVM option or import the certificate to any directory.</p> <p>For the serverless runtime environment, specify the keystore file path in the serverless agent directory.</p> <p>For example, <code>/home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</code></p>
KeyStore Password	<p>The password for the keystore file required for secure communication.</p> <p>You can also configure the keystore password as a JVM option.</p>
Proxy Type	<p>Type of proxy.</p> <p>You can select one of the following options:</p> <ul style="list-style-type: none"> - No Proxy: Bypasses the proxy server configured in the agent or the connection properties. - Platform Proxy: Considers the proxy configured in the agent. - Custom Proxy: Considers the proxy configured in the connection properties.
Proxy Configuration	<p>The format required to configure proxy.</p> <p>You can configure proxy using the following format: <code><host>:<port></code></p> <p>You cannot configure an authenticated proxy server.</p>
Advanced Fields	<p>Enter the arguments that the agent uses when connecting to a REST endpoint.</p> <p>You can specify the following arguments, each separated by a semicolon (;):</p> <ul style="list-style-type: none"> - ConnectionTimeout. The wait time in milliseconds to get a response from a REST endpoint. The connection ends after the connection timeout is over. Default is the timeout defined in the endpoint API. - Note: If you define both the REST V2 connection timeout and the endpoint API timeout, the connection ends at the shortest defined timeout. - connectiondelaytime. The delay time in milliseconds to send a request to a REST endpoint. Default is 10000. - retryattempts. Number of times the connection is attempted when 400 and 500 series error codes are returned in the response. Default is 3. Specify 0 to disable the retry attempts. - qualifiedSchema. Determines if the schema selected is qualified or unqualified. Default is false. <p>For example,</p> <pre>connectiondelaytime:10000;retryattempts:5</pre> <p>Note: In a streaming ingestion and replication task, only <code>ConnectionTimeout</code> and <code>retryattempts</code> are applicable.</p>

Table 2. ForOAuth 2.0-Authorization Code authentication

Connection property	Description
Authorization Token URL	Authorization server URL configured in your application.
Access Token URL	Access token URL configured in your application.

Connection property	Description
Client ID	Client ID of your application.
Client Secret	Client secret of your application.
Scope	Specifies access control if the API endpoint has defined custom scopes. Enter space separated scope attributes. For example, root_readonly root_readwrite manage_app_users
Access Token Parameters	Additional parameters to use with the access token URL. Parameters must be defined in the JSON format. For example, [{"Name": "resource", "Value": "https://<serverName>"}]
Authorization Code Parameters	Additional parameters to use with the authorization token URL. Parameters must be defined in the JSON format. For example, [{"Name": "max_age", "Value": 60}, {"Name": "state", "Value": "test"}]
Client Authentication	Select an option to send Client ID and Client Secret for authorization either in the request body or in the request header. Default is Send Client Credentials in Body .
Generate Access Token	Generates access token and refresh token based on the information provided in the above fields.
Access Token	Enter the access token value or click Generate Access Token to populate the access token value. To pass the generate access token call through a proxy server, you must configure an unauthenticated proxy server at the Secure Agent level. The REST V2 connection-level proxy configuration does not apply to the generate access token call.
Refresh Token	Enter the refresh token value or click Generate Access Token to populate the refresh token value. If the access token is not valid or expires, the Secure Agent fetches a new access token with the help of refresh token. If the refresh token expires, you must either provide a valid refresh token or regenerate a new refresh token by clicking Generate Access Token .
Swagger File Path	The path of the Swagger file or OpenAPI file. You can specify one of the following file paths: - Absolute path along with the file name - Hosted URL If you provide the absolute path of the swagger file or OpenAPI file, the file must be located on the Secure Agent machine. The hosted URL must return the content of the file without prompting for further authentication and redirection. For example, the path of the swagger file can be: C:\swagger\sampleSwagger.json The user must have the read permission for the folder and the file. Note: In a streaming ingestion and replication task, use only a hosted URL of the swagger specification file as the swagger file path.

Connection property	Description
TrustStore File Path	<p>The absolute path of the truststore file that contains the TLS certificate to establish a one-way or two-way secure connection with the REST API. Specify a directory path that is available on each Secure Agent machine.</p> <p>You can also configure the truststore file name and password as a JVM option or import the certificate to the following directory:</p> <p><Secure Agent installation directory>\jre\lib\security\cacerts.</p> <p>For the serverless runtime environment, specify the truststore file path in the serverless agent directory.</p> <p>For example, /home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</p>
TrustStore Password	<p>The password for the truststore file that contains the SSL certificate.</p> <p>You can also configure the truststore password as a JVM option.</p>
KeyStore File Path	<p>The absolute path of the keystore file that contains the keys and certificates required to establish a two-way secure communication with the REST API. Specify a directory path that is available on each Secure Agent machine.</p> <p>You can also configure the keystore file name and location as a JVM option or import the certificate to any directory.</p> <p>For the serverless runtime environment, specify the keystore file path in the serverless agent directory.</p> <p>For example, /home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</p>
KeyStore Password	<p>The password for the keystore file required for secure communication.</p> <p>You can also configure the keystore password as a JVM option.</p>
Proxy Type	<p>Type of proxy. You can select one of the following options:</p> <ul style="list-style-type: none"> - No Proxy: Bypasses the proxy server configured in the agent or the connection properties. - Platform Proxy: Considers the proxy configured in the agent. - Custom Proxy: Considers the proxy configured in the connection properties.

Connection property	Description
Proxy Configuration	<p>The format required to configure proxy.</p> <p>You can configure proxy using the following format: <host>:<port></p> <p>You cannot configure an authenticated proxy server.</p>
Advanced Fields	<p>Enter the arguments that the agent uses when connecting to a REST endpoint.</p> <p>You can specify the following arguments, each separated by a semicolon (;):</p> <ul style="list-style-type: none"> - ConnectionTimeout. The wait time in milliseconds to get a response from a REST endpoint. The connection ends after the connection timeout is over. Default is the timeout defined in the endpoint API. - Note: If you define both the REST V2 connection timeout and the endpoint API timeout, the connection ends at the shortest defined timeout. - connectiondelaytime. The delay time in milliseconds to send a request to a REST endpoint. Default is 10000. - retryattempts. Number of times the connection is attempted when 400 and 500 series error codes are returned in the response. Default is 3. Specify 0 to disable the retry attempts. - qualifiedSchema. Determines if the schema selected is qualified or unqualified. Default is false. <p>For example,</p> <pre>connectiondelaytime:10000;retryattempts:5</pre> <p>Note: In a streaming ingestion and replication task, only ConnectionTimeout and retryattempts are applicable.</p>

Table 3. For JWT Bearer Token authentication

Connection property	Description
JWT Header	<p>JWT header in JSON format.</p> <p>Sample:</p> <pre>{ "alg": "RS256", "kid": "xyyyz" }</pre> <p>You can configure HS256 and RS256 algorithms.</p>
JWT Payload	<p>JWT payload in JSON format.</p> <p>Sample:</p> <pre>{ "iss": "abc", "sub": "678", "aud": "https://api.box.com/oauth2/token", "box_sub_type": "enterprise", "exp": "120", "jti": "3ee9364e" }</pre> <p>The expiry time represented as exp is the relative time in seconds. The expiry time is calculated in the UTC format from the token issuer time (iat).</p> <p>When iat is defined in the payload and the expiry time is reached, mappings and Generate Access Token will fail.</p>

Connection property	Description
	<p>To generate a new access token, you must provide a valid <code>iat</code> in the payload.</p> <p>If <code>iat</code> is not defined in the payload, the expiry time is calculated from the current timestamp.</p> <p>To pass the expiry time as a string value, enclose the value with double quotes. For example:</p> <pre>"exp": "120",</pre> <p>To pass the expiry time as an integer value, do not enclose the value with double quotes.</p> <p>For example,</p> <pre>"exp": 120,</pre>
Authorization Server	Access token URL configured in your application.
Authorization Advanced Properties	<p>Additional parameters to use with the access token URL. Parameters must be defined in the JSON format.</p> <p>For example,</p> <pre>[{"Name": "client_id", "Value": "abc"}, \ {"Name": "client_secret", "Value": "abc"}]</pre>
TrustStore File Path	<p>The absolute path of the truststore file that contains the TLS certificate to establish a one-way or two-way secure connection with the REST API. Specify a directory path that is available on each Secure Agent machine.</p> <p>You can also configure the truststore file name and password as a JVM option or import the certificate to the following directory:</p> <pre><Secure Agent installation directory>\jre \lib\security\cacerts.</pre> <p>For the serverless runtime environment, specify the truststore file path in the serverless agent directory.</p> <p>For example, <code>/home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</code></p>
TrustStore Password	<p>The password for the truststore file that contains the SSL certificate.</p> <p>You can also configure the truststore password as a JVM option.</p>
KeyStore File Path	<p>Mandatory. The absolute path of the keystore file that contains the keys and certificates required to establish a two-way secure communication with the REST API. Specify a directory path that is available on each Secure Agent machine.</p> <p>You can also configure the keystore file name and location as a JVM option or import the certificate to any directory.</p> <p>For the serverless runtime environment, specify the keystore file path in the serverless agent directory.</p> <p>For example, <code>/home/cldagnt/SystemAgent/serverless/configurations/ssl_store/<cert_name>.jks</code></p>

Connection property	Description
KeyStore Password	<p>Mandatory. The password for the keystore file required for secure communication.</p> <p>You can also configure the keystore password as a JVM option.</p>
Private Key Alias	<p>Mandatory. Alias name of the private key used to sign the JWT payload.</p>
Private Key Password	<p>Mandatory. The password for the keystore file required for secure communication. The private key password must be same as the keystore password.</p>
Access Token	<p>Enter the access token value or click Generate Access Token to populate the access token value.</p> <p>To pass the generate access token call through a proxy server, you must configure an unauthenticated proxy server at the Secure Agent level. The REST V2 connection-level proxy configuration does not apply to the generate access token call.</p>
Swagger File Path	<p>The path of the Swagger file or OpenAPI file.</p> <p>You can specify one of the following file paths:</p> <ul style="list-style-type: none"> - Absolute path along with the file name - Hosted URL <p>If you provide the absolute path of the swagger file or OpenAPI file, the file must be located on the Secure Agent machine.</p> <p>The hosted URL must return the content of the file without prompting for further authentication and redirection.</p> <p>For example, the path of the swagger file can be:</p> <pre>C:\swagger\sampleSwagger.json</pre> <p>The user must have the read permission for the folder and the file.</p> <p>Note: In a streaming ingestion and replication task, use only a hosted URL of the swagger specification file as the swagger file path.</p>
Proxy Type	<p>Type of proxy. You can select one of the following options:</p> <ul style="list-style-type: none"> - No Proxy. Bypasses the proxy server configured in the agent or the connection properties. - Platform Proxy. Considers the proxy configured in the agent. - Custom Proxy. Considers the proxy configured in the connection properties.

Connection property	Description
Proxy Configuration	<p>The format required to configure proxy. You can configure proxy using the following format: <host>:<port></p> <p>You cannot configure an authenticated proxy server.</p>
Advanced Fields	<p>Enter the arguments that the agent uses when connecting to a REST endpoint. You can specify the following arguments, each separated by a semicolon (;):</p> <ul style="list-style-type: none"> - ConnectionTimeout. The wait time in milliseconds to get a response from a REST endpoint. The connection ends after the connection timeout is over. Default is the timeout defined in the endpoint API. - Note: If you define both the REST V2 connection timeout and the endpoint API timeout, the connection ends at the shortest defined timeout. - connectiondelaytime. The delay time in milliseconds to send a request to a REST endpoint. Default is 10000. - retryattempts. Number of times the connection is attempted when 400 and 500 series error codes are returned in the response. Default is 3. Specify 0 to disable the retry attempts. - qualifiedSchema. Determines if the schema selected is qualified or unqualified. Default is false. <p>For example, connectiondelaytime:10000;retryattempts:5</p> <p>Note: In a streaming ingestion and replication task, only ConnectionTimeout and retryattempts are applicable.</p>

Salesforce Marketing Cloud connection properties

When you set up a Salesforce Marketing Cloud connection, configure the connection properties.

The following table describes the Salesforce Marketing Cloud connection properties:

Property	Description
Connection Name	<p>Name of the connection.</p> <p>Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.</p>
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Salesforce Marketing Cloud connection type.

Property	Description
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run an application ingestion and replication task on a Hosted Agent or serverless runtime environment.
Salesforce Marketing Cloud Url	The URL that the agent uses to connect to the Salesforce Marketing Cloud WSDL.
Username	Applies to basic authentication. The user name of the Salesforce Marketing Cloud account. Note: This property is not applicable to connections configured for application ingestion and replication tasks.
Password	Applies to basic authentication. The password for the Salesforce Marketing Cloud account. Note: This property is not applicable to connections configured for application ingestion and replication tasks.
Client ID	The client ID of Salesforce Marketing Cloud required to generate a valid access token.
Client Secret	The client secret of Salesforce Marketing Cloud required to generate a valid access token.
Use Proxy Server	Connects to Salesforce Marketing Cloud through proxy. Note: This property is not applicable to connections configured for application ingestion and replication tasks.
Enable Logging	Enables logging for the task. When you enable logging, you can view the session log for the log details. Note: This property is not applicable to connections configured for application ingestion and replication tasks.
UTC offset	Uses the UTC offset connection property to read data from and write data to Salesforce Marketing Cloud in the UTC offset time zone. Note: This property is not applicable to connections configured for application ingestion and replication tasks.
Batch Size	Number of rows that the agent writes in a batch to the target. When you insert or update data and specify the contact key, the data associated with the specified contact ID is inserted or updated in a batch to Salesforce Marketing Cloud. When you upsert data to Salesforce Marketing Cloud, do not specify the contact key. Note: This property is not applicable to connections configured for application ingestion and replication tasks.
Enable Multiple BU	Uses the Salesforce Marketing Cloud connection to access data across all business units. Select this option if there are multiple business units in your Salesforce Marketing Cloud account. Note: This property is not applicable to connections configured for application ingestion and replication tasks.

Salesforce Mass Ingestion connection properties

When you set up a Salesforce Mass Ingestion connection, you must configure the connection properties.

The Salesforce Mass Ingestion connection uses a connected app to access the Salesforce data. Before you configure the connection, you must configure a connected app in Salesforce to allow the connection to access the Salesforce data.

Note: For more information about configuring a connected app, see the Knowledge Base article [000172095](#).

The properties of a Salesforce Mass Ingestion connection vary based on the authentication method that you specify for the connection. When you create a connection, you can select one of the following authentication methods:

- **OAuth 2.0 Username-Password Flow:** Authenticates the connection by using the Salesforce account login credentials and the consumer key and consumer secret that Salesforce generates for the connected app.
- **OAuth 2.0 JWT Bearer Flow:** Authenticates the connection by using the Salesforce account user name, private key alias, private key password, and the consumer key that Salesforce generates for the connected app. Informatica recommends that you use this authentication method because this method provides secured access to Salesforce without sharing sensitive information, such as consumer secret and Salesforce account password.

Connection properties for OAuth 2.0 Username-Password Flow authentication

The following table describes the connection properties for a Salesforce Mass Ingestion connection configured with OAuth 2.0 Username-Password Flow authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Salesforce Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
User Name	User name of the Salesforce account.
Password	Password for the Salesforce account.
Security Token	Security token associated with the Salesforce account. You can configure the connection without specifying the security token if there are no IP restrictions specified for the connected app. However, you must specify the security token if IP restrictions are enforced for the connected app and if the Secure Agent is not running on the trusted IP range specified for your Salesforce organization. Note: If you do not have the security token, reset the security token in Salesforce. For more information about resetting the security token, see the Salesforce documentation .

Connection property	Description
Consumer Key	Consumer key that Salesforce generates when you enable OAuth 2.0 authentication for the connected app.
Consumer Secret	Consumer secret that Salesforce generates when you enable OAuth 2.0 authentication for the connected app.
API Version	Version of the Salesforce API that you want to use to access the source data. Default is 51.0. Note: You cannot use a version older than 51.0.
OAuth token URL	OAuth 2.0 token endpoint of the Salesforce organization. The connected app sends access token requests to this endpoint. Default value is: <code>https://login.salesforce.com/services/oauth2/token</code> This default URL is used for all Salesforce instances. Alternatively, you can enter an instance-specific URL: <code>https://<instance domain URL>/services/oauth2/token</code> An instance-specific URL can establish a more direct and faster connection to the Salesforce host server. If the load on the common default endpoint is heavy and ingestion jobs fail with an authentication error when using it, use this alternative URL instead.

Note: For more information about the OAuth 2.0 Username-Password Flow authentication method, see the Salesforce documentation.

Connection properties for OAuth 2.0 JWT Bearer Flow authentication

The following table describes the connection properties for a Salesforce Mass Ingestion connection configured with OAuth 2.0 JWT Bearer Flow authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Salesforce Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
User Name	User name of the Salesforce account.
Consumer Key	Consumer key that Salesforce generates when you enable OAuth 2.0 authentication for the connected app.

Connection property	Description
Keystore Path	Absolute path to the keystore file that contains the X509 certificate required to validate a JSON Web Token (JWT) and establish a secure connection with Salesforce. The keystore file must be in the Java KeyStore (JKS) format.
Keystore Password	Password for the keystore file.
Private Key Alias	Alias name of the private key used to sign the JWT.
Private Key Password	Password for the private key.
API Version	Version of the Salesforce API that you want to use to access the source data. Default is 51.0. Note: You cannot use a version older than 51.0.
OAuth token URL	OAuth 2.0 token endpoint of the Salesforce organization. The connected app sends access token requests to this endpoint. Default value is: <code>https://login.salesforce.com/services/oauth2/token</code> This default URL is used for all Salesforce instances. Alternatively, you can enter an instance-specific URL: <code>https://<instance domain URL>/services/oauth2/token</code> An instance-specific URL can establish a more direct and faster connection to the Salesforce host server. If the load on the common default endpoint is heavy and ingestion jobs fail with an authentication error when using it, use this alternative URL instead.

Note: For more information about the OAuth 2.0 JWT Bearer Flow authentication method, see the Salesforce documentation.

SAP HANA Database Ingestion connection properties

When you set up an SAP HANA connection for a database ingestion and replication task, you must configure connection properties.

The following table describes the SAP HANA connection properties:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.

Connection property	Description
Type	Select SAP HANA Database Ingestion as the connection type.
Runtime Environment	The name of the runtime environment where you want to run database ingestion and replication tasks. You define runtime environments in Administrator.
User Name	The user name for connecting to the SAP HANA instance. Enter the user name in the same case as in the database user name specified in SAP HANA.
Password	The password for connecting to the SAP HANA instance.
Host	The name of the machine that hosts the SAP HANA database server.
Port	The port number for the SAP HANA server to which you want to connect. Default is 30015.
Database Name	The SAP HANA source database name.
Advanced Connection Properties	Optional advanced properties for the SAP HANA JDBC driver, which is used to connect to the SAP HANA source. If you specify more than one <i>property=value</i> entry, separate them with an ampersand (&). The JDBC connection properties that you can enter in this field are described in the SAP JDBC Connection Properties documentation. For example: encrypt=true.
Capture Type	<p>Select one of the following options to indicate the capture method that database ingestion incremental load jobs use to capture change data from SAP HANA databases:</p> <ul style="list-style-type: none"> - Trigger Based. Capture change data from SAP HANA source tables in the schema by using AFTER DELETE, AFTER INSERT, and AFTER UPDATE triggers. The triggers get before images and after images of DML changes for each source table and write entries for the changes to the PKLOG and shadow _CDC tables. This method is the original capture method. - Log Based (Preview). Capture change data from the SAP HANA database logs. This method is available only in Preview mode. Preview functionality is supported for evaluation purposes but is unwarranted and is not supported in production environments or any environment that you plan to push to production. For more information, contact Informatica Global Customer Support.
Log Clear	<p>Required for incremental loads. Enter the time interval, in days, after which the PKLOG table entries and shadow _CDC table entries are purged. The purging occurs only while an incremental load job is running.</p> <p>Valid values for a database ingestion job are 0 to 366. Any positive value in this range cause automatic housekeeping to run while the incremental job is running. Default is 14.</p> <p>A value of 0 means that the table entries are not purged. For manual housekeeping, enter 0 and use your in-house process.</p> <p>Any value outside the range of 0 to 366, including a negative number or non-numeric value, causes database ingestion jobs that use the connection to fail with the following error:</p> <pre>LogClear contains a non numeric number. Caused by: LogClear contains a non numeric number.</pre>
Trigger Prefix	If you use the Trigger Based capture type, you can add a prefix to the names of the AFTER DELETE, AFTER INSERT, and AFTER UPDATE triggers that the CDC script generates for each source table to get before images and after images of the DML changes. Enter any prefix value up to 16 characters in length. An underscore (_) follows the prefix in the trigger name, for example, TX_SAP_DEMO_TABLE_DBMI_USER_t_d . You can use the prefix to comply with your site's trigger naming conventions.
Cache Type	If you selected the Log Based (Preview) capture type, select Hana or Oracle as the cache type.

Connection property	Description
Cache Host	If you selected the Log Based (Preview) capture type, enter the host name of the machine that hosts the cache database.
Cache Port	If you selected the Log Based (Preview) capture type, enter the port number for the cache database server.
Cache User Name	If you selected the Log Based (Preview) capture type, enter the user name to use for connecting to the cache database.
Cache Password	If you selected the Log Based (Preview) capture type, enter the password to use for connecting to the cache database.
Cache Database/ Service Name	If you selected the Log Based (Preview) capture type, enter either the Hana cache database name or the Oracle cache service name, depending on the cache type you selected.
Cache Additional Connection Properties	<p>If you selected the Log Based (Preview) capture type, you can enter a list of optional cache connection properties. If you use Hana cache, use the ampersand (&) separator. If you use Oracle cache, use the semicolon (;) separator.</p> <p>Examples:</p> <p>Hana: latency=0&communicationtimeout=0</p> <p>Oracle: EncryptionMethod=SSL;CryptoProtocolVersion=TLSv1.1</p>
Cache Security Connection Properties	<p>If you selected the Log Based (Preview) capture type, you can enter a list of optional security properties for the cache connection. If you use Hana cache, use the ampersand (&) separator. If you use Oracle cache, use the semicolon (;) separator.</p> <p>Examples:</p> <p>Hana: encrypt=true&validateCertificate=false</p> <p>Oracle: KeyStorePassword=xyz;TrustStorePassword=xy</p>
Server Log Path	If you selected the Log Based (Preview) capture type, enter the log path for the SAP HANA DB server.
Client Log Path	If you selected the Log Based (Preview) capture type, enter the mapping of the Secure Agent machine mount path to the source database log location.
Client Archive Log Path	If you selected the Log Based (Preview) capture type, enter the mapping of the Secure Agent machine mount path to the source database archive log location.

Note: If you test the connection and the test fails, check that the SAP HANA JDBC driver file, ngdbc.jar, has been installed at *Secure Agent installation directory>/ext/connectors/thirdparty/informatica.hanami*.

SAP Mass Ingestion connection properties

When you set up a SAP Mass Ingestion connection, you must configure the connection properties.

The following table describes the connection properties for a SAP Mass Ingestion connection:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be SAP Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
User Name	User name of the SAP instance.
Password	Password for the SAP instance.
Language Code	Language code that corresponds to the SAP language.
System Number	System number of the SAP server.
Client Number	Client number of the SAP server.
Port Range	HTTP port range to run the Netty server.
Connection Type	Type of connection to access the ABAP application server. Options are: <ul style="list-style-type: none">- Direct Connection: Accesses a single ABAP application server using the server host.- Load Balancing Connection: Accesses a group of ABAP application servers through the message server.
Application Server	Name of the SAP application server host. Note: This field appears only for the Direct Connection type.
Message Server	IP address or name of the SAP message server. Note: This field appears only for the Load Balancing Connection type.
SAP Logon Group	Name of the group of servers that belong to the SAP system you want to access. Note: This field appears only for the Load Balancing Connection type.
SAP System ID	ID of the SAP system that you want to access. Note: This field appears only for the Load Balancing Connection type.
Message Server Port	Port number on which the SAP message server is listening. Note: This field appears only for the Load Balancing Connection type.

Connection property	Description
Database	The name of the underlying database. Select one of the following options: <ul style="list-style-type: none"> - Oracle - SAP HANA (S/4 trigger based)
For Oracle database	
Database user name	User name of the database instance.
Database password	Password for the database instance.
Host	Host name of the database server.
Port	Network port number used to connect to the database server. Default is 1521.
Service Name	Service name or System ID (SID) that uniquely identifies the Oracle database. Specify the SID in the following format to connect to Oracle databases: SID:<ORACLE_SID>
Code Page	The code page of the database server. Application ingestion and replication tasks use the UTF-8 code page. Default is UTF-8.
Encryption Method	For initial load jobs, determines whether the data exchanged between the Secure Agent and the Oracle database server is encrypted: Select one of the following options: <ul style="list-style-type: none"> - SSL. Establishes a secure connection using SSL for data encryption. If the Oracle database server cannot configure SSL, the connection fails. - No Encryption. Establishes a connection without using SSL. Data is not encrypted. Default is No Encryption .
Crypto Protocol Version	If you selected SSL as the encryption method, you must specify a cryptographic protocol or a list of cryptographic protocols supported by your server to use with an encrypted connection. Select one of the following options: <ul style="list-style-type: none"> - SSLv2 - SSLv3 - TLSv1.2 Default is TLSv1.2 .
Validate Server Certificate	If you selected SSL as the encryption method, this option controls whether the Secure Agent validates the server certificate that is sent by the Oracle database server. Select one of the following options: <ul style="list-style-type: none"> - True. Validate the server certificate. - False. Do not validate the server certificate. Default is False . If you also specify the Host Name in Certificate property, the Secure Agent also validates the host name in the certificate.
Trust Store	If you selected SSL as the encryption method and enabled validation of the server certificate, specify the path and name of the truststore file, which contains the list of the Certificate Authorities (CAs) that the client trusts for SSL authentication.

Connection property	Description
Trust Store Password	If you selected SSL as the encryption method and enabled validation of the server certificate, specify a password for accessing the contents of the truststore file.
Host Name in Certificate	If you selected SSL as the encryption method and enabled validation of the server certificate, specify the host name of the machine that hosts the Oracle database to provide for additional security. The Secure Agent validates the host name included in the connection with the host name in the SSL certificate.
Key Store	If you selected SSL as the encryption method and client authentication is enabled on the Oracle database server, specify the path and name of the keystore file. The keystore file contains the certificates that the client sends to the Oracle server in response to the server's certificate request.
Key Store Password	If you selected SSL as the encryption method and client authentication is enabled on the Oracle database server, specify the password for the keystore file.
Key Password	If you selected SSL as the encryption method and client authentication is enabled on the Oracle database server, specify the password for the keys in the keystore file. Use this property when the keys have a different password than the keystore file.
Database Connect String	An Oracle connection string, defined in TNS, that application ingestion and replication tasks use to connect to the Oracle database.
TDE Wallet Directory	<p>The path to the directory that contains the Oracle wallet file used for Oracle Transparent Data Encryption (TDE). Specify this property value only if you capture change data from TDE-encrypted table spaces and one of the following conditions are true:</p> <ul style="list-style-type: none"> - The Oracle wallet is not available to the database. - The Oracle database is running on a server that is remote from Oracle redo logs. - The wallet directory is not in the default location on the database host or the wallet name is not the default name of ewallet.p12. - The wallet directory is not available to the Secure Agent host.
TDE Wallet Password	A clear text password that is required to access the Oracle TDE wallet and get the master key. This property value is required if you need to read and decrypt data from TDE-encrypted tablespaces in the Oracle source database.
Directory Substitution	<p>A local path prefix to substitute for the server path prefix of the redo logs on the Oracle server. This substitute local path is required when the log reader runs on a system other than the Oracle server and uses a different mapping to access the redo log files.</p> <p>Use this property in the following situations:</p> <ul style="list-style-type: none"> - The redo logs reside on shared disk. - The redo logs have been copied to a system other than the Oracle system. - The archived redo logs are accessed by using a different NFS mount. <p>Note: Do not use this property if you use Oracle Automatic Storage Management (ASM) to manage the redo logs.</p> <p>You can define one or more substitutions. Use the following format:</p> <pre>server_path_prefix,local_path_prefix;server_path_prefix,local_path_prefix;...</pre>

Connection property	Description
Reader Active Log Mask	<p>A mask that the log reader uses for selecting active redo logs when the Oracle database uses multiplexing of redo logs. The log reader compares the mask against the member names in an active redo log group to determine which log to read. In the mask, you can use the asterisk (*) wildcard to represent zero or more characters.</p> <p>The mask can be up to 128 characters in length. It is case-sensitive on Linux or UNIX systems but not on Windows systems.</p>
Reader Archive Destination 1	<p>The primary log destination from which the log reader reads archived logs, when Oracle is configured to write more than one copy of each archived redo log. Enter a number that corresponds to an <i>n</i> value in an Oracle LOG_ARCHIVE_DEST_<i>n</i> initialization parameter, where <i>n</i> is a value from 1 to 10. If you set only one of the Reader Archive Destination 1 and Destination 2 properties, the log reader uses that property setting. If you specify neither property, the archive log queries are not filtered by the log destination.</p>
Reader Archive Destination 2	<p>The secondary log destination from which the log reader reads archived logs when the primary destination becomes unavailable or when the logs at the primary destination cannot be read. For example, logs might have been corrupted or deleted. Enter a number that corresponds to the <i>n</i> value in an Oracle LOG_ARCHIVE_DEST_<i>n</i> initialization parameter, where <i>n</i> is a value from 1 to 10. Usually, this value is a number greater than 1.</p>
Reader ASM Connect String	<p>In an Oracle ASM environment, the Oracle connection string, defined in TNS, that the log reader uses to connect to the ASM instance that manages storage of active and archived redo logs for the source database.</p>
Reader ASM User Name	<p>In an Oracle ASM environment, an Oracle user ID that the log reader uses to connect to the ASM instance that manages storage of active and archived redo logs for the source database. This user ID must have SYSDBA or SYSASM authority. To use SYSASM authority, set the Reader ASM Connect As SYSASM property to Y.</p>
Reader ASM Password	<p>In an Oracle ASM environment, a clear text password for the user that is specified in the Reader ASM User Name property. The log reader uses this password and the ASM user name to connect to the ASM instance that manages storage of active and archived redo logs for the source database.</p>
Reader ASM Connect As SYSASM	<p>If you use Oracle 11g ASM or later and want the log reader to use a user ID that has SYSASM authority to connect to the ASM instance, select this check box. Also specify a user ID that has SYSASM authority in the Reader ASM User Name property. To use a user ID that has SYSDBA authority, clear this check box. By default, this check box is cleared.</p>

Connection property	Description
Reader Mode	<p>Indicates the source of and types of Oracle redo logs that the log reader reads. Select one of the following options:</p> <ul style="list-style-type: none"> - ACTIVE. Read active and archived redo logs from the Oracle online system. Optionally, you can use the Reader Active Log Mask property to filter the active redo logs and use the Reader Archive Destination 1 and Reader Archive Destination 2 properties to limit the archived log destinations from which to read archived logs. - ARCHIVEONLY. Read only archived redo logs. Optionally, you can use the Reader Archive Destination 1 and Reader Archive Destination 2 properties to limit the archived log destinations from which to read archived logs. - ARCHIVECOPY. Read archived redo logs that have been copied to an alternate file system. Use this option in the following situations: <ul style="list-style-type: none"> - You do not have the authority to access the Oracle archived redo logs directly. - The archived redo logs are written to ASM, but you do not have access to ASM. - The archived log retention policy for the database server causes the archived logs to not be retained long enough. <p>With this option, the Reader Archive Destination 1 and Reader Archive Destination 2 properties are ignored.</p> <p>Default is ACTIVE.</p>
Reader Standby Log Mask	<p>A mask that the log reader uses for selecting redo logs for an Oracle physical standby database when the database uses multiplexing of redo logs. The log reader compares the mask against the member names in an redo log group to determine which log to read. In the mask, you can use the asterisk (*) wildcard to represent zero or more characters.</p> <p>The mask can be up to 128 characters in length. It is case-sensitive on Linux or UNIX systems but not on Windows systems.</p>
Standby Connect String	<p>An Oracle connection string, defined in TNS, that the log reader uses to connect to the Oracle physical standby database for change capture when the database is not open with read only access.</p>
Standby User Name	<p>A user ID that the log reader uses to connect to the Oracle physical standby database for change capture. This user ID must have SYSDBA authority.</p>
Standby Password	<p>A password that the log reader uses to connect to the Oracle physical standby database for change capture.</p>
RAC Members	<p>The maximum number of active redo log threads, or <i>members</i>, in an Oracle Real Application Cluster (RAC) that can be tracked. For a Data Guard physical standby database that supports a primary database in a RAC environment, this value is the number of active threads for the primary database.</p> <p>Valid values are 1 to 100. Default is 0, which causes an appropriate number of log threads to be determined automatically. If this value is not appropriate for your environment, set this property to a value greater than 0.</p>

Connection property	Description
BFILE Access	<p>Select this check box in the following circumstances:</p> <ul style="list-style-type: none"> - You use BFILE access to redo logs in physical directories on the local Oracle server file system. BFILE access uses Oracle directory objects to remotely access the redo logs in the file system. This method is an alternative to other log access methods such as ASM or NFS mounts. - You have an Amazon Relational Database Service (RDS) for Oracle source. In this case, this option enables access to the redo logs of a cloud-based database instance deployed in RDS. <p>By default, this check box is cleared.</p>
For SAP HANA (S/4 trigger based) database	
User Name	The user name to connect to the SAP HANA instance.
Password	The password to connect to the SAP HANA instance.
Host	The name of the machine that hosts the SAP HANA database server.
Port	The port number of the SAP HANA server that you want to connect to. Default is 30015.
Database Name	The SAP HANA source database name.
Advanced Connection Properties	<p>Advanced properties for the SAP HANA JDBC driver, which is used to connect to the SAP HANA source. If you specify more than one <i>property=value</i> entry, separate them with an ampersand (&). The JDBC connection properties that you can enter in this field are described in the SAP JDBC Connection Properties documentation. For example: encrypt=true.</p>
Log Clear	<p>Required for incremental loads. The time interval, in days, after which the PKLOG table entries and shadow _CDC table entries are purged. The purging occurs only while an incremental load job is running.</p> <p>Valid values for a database ingestion job are 0 to 366. Any positive value in this range causes automatic housekeeping to run while the incremental job is running. Default is 14.</p> <p>A value of 0 means that the table entries are not purged. For manual housekeeping, enter 0 and use your in-house process.</p> <p>Any value outside the range of 0 to 366, including a negative number or non-numeric value, causes database ingestion jobs that use the connection to fail with the following error:</p> <pre>LogClear contains a non numeric number. Caused by: LogClear contains a non numeric number.</pre>
Trigger Prefix	<p>If you use the trigger-based capture type, you can add a prefix to the names of the AFTER DELETE, AFTER INSERT, and AFTER UPDATE triggers that the CDC script generates for each source table to get before images and after images of the DML changes. Enter any prefix value up to 16 characters in length. An underscore (_) follows the prefix in the trigger name, for example, TX_SAP_DEMO_TABLE_DBML_USER_t_d. You can use the prefix to comply with your site's trigger naming conventions.</p>

SAP ODP Extractor connection properties

When you set up an **SAP ODP Extractor** connection, configure the connection properties.

The following table describes the SAP ODP Extractor connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	SAP ODP Extractor
Runtime Environment	The name of the runtime environment where you want to run tasks to access SAP S/4HANA or SAP ECC.
SAP Server Connection Type	The SAP server connection type to use. Select from the following options: <ul style="list-style-type: none">- Application Server Connection. Connect to an SAP Application Server using the SAP user name and password.- Application Server SNC Connection. Connect to an SAP Application Server using the secured network connection:<ul style="list-style-type: none">- With X.509 Certificate. You do not need to specify the SAP user name and password explicitly. You must provide the path of the x.509 certificate file.- Without X.509 Certificate. You must provide the SAP user name.- Load Balancing Server Connection. Connect to an SAP Application Server with the least load at run time.- Load Balancing Server SNC Connection. Connect to an SAP Application Server using SNC with the least load at run time. Note: Before you use an SNC connection, you must verify that SNC is configured both on the SAP Server and the machine where the Secure Agent runs.

The following table describes the properties that must configure when you select **Application Server Connection** as the connection type:

Connection property	Description
SAP Client Number	The client number of the SAP Server.
SAP Language	Language code that corresponds to the SAP language.
SAP Application Server	The host name of the SAP Application Server.
SAP System Number	The system number of the SAP Server to connect.
SAP Username	The SAP user name with the appropriate user authorization.
SAP Password	The SAP password.

Connection property	Description
Subscriber Name	A name which defines the Secure Agent as a unique subscriber in the SAP system. SAP uses this name to define unique operational delta queue (ODQ) in case of delta read from ODP.
Additional Parameters	<p>Additional SAP parameters that the Secure Agent uses to connect to the SAP system. For example, to generate SAP JCo and SAP CPIC trace, specify the following properties:</p> <pre>jco.client.trace="1"; jco.client.cpic_trace="3";</pre> <p>During the runtime, the JCo and CPIC traces file are generated in the following location: <Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>\ICS\main\bin\rdtm</p> <p>During the design time, the CPIC traces are generated in the tomcat.out files at the following location: <Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>tomcat.out</p>
Display Delta Fields	<p>Specifies whether the mapping displays the operation modes that caused the changed data on ODP sources.</p> <p>When enabled, the mapping generates the ODQ_CHANGEMODE and ODQ_ENTITYCNTR fields on the Fields tab for ODP sources that are enabled with Operational Delta Queue (ODQ).</p> <p>Default is disabled.</p>

The following table describes the properties that must configure when you select **Load Balancing Server Connection** as the connection type:

Connection property	Description
SAP Client Number	The client number of the SAP Server.
SAP Language	Language code that corresponds to the SAP language.
SAP Message Server	Host name of the SAP Message Server.
SAP System ID	The system ID of the SAP Message Server.
SAP Group	The login group name, for example, PUBLIC.
SAP Username	The SAP user name with the appropriate user authorization.
SAP Password	The SAP password.
Subscriber Name	<p>A name which defines the Secure Agent as a unique subscriber in the SAP system. SAP uses this name to define unique operational delta queue (ODQ) in case of delta read from ODP.</p>

Connection property	Description
Additional Parameters	<p>Additional SAP parameters that the Secure Agent uses to connect to the SAP system. For example, to generate SAP JCo and SAP CPIC trace, specify the following properties:</p> <pre>jco.client.trace="1"; jco.client.cpic_trace="3";</pre> <p>During the runtime, the JCo and CPIC traces file are generated in the following location: <Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>\ICS\main\bin\rdtm</p> <p>During the design time, the CPIC traces are generated in the tomcat.out files at the following location: <Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>tomcat.out</p>
Display Delta Fields	<p>Specifies whether the mapping displays the operation modes that caused the changed data on ODP sources.</p> <p>When enabled, the mapping generates the ODQ_CHANGEMODE and ODQ_ENTITYCNTR fields on the Fields tab for ODP sources that are enabled with Operational Delta Queue (ODQ). Default is disabled.</p>

The following table describes the properties that must configure when you select **Application Server SNC Connection** as the connection type:

Connection property	Description
SAP Client Number	The client number of the SAP Server.
SAP Language	Language code that corresponds to the SAP language.
SAP Application Server	The host name of the SAP Application Server.
SAP System Number	The system number of the SAP Server to connect.
SNC My Name	Optional. The Informatica client Personal Security Environment (PSE) or certificate name. Default length is 256.
SNC Partner Name	The Informatica client PSE or certificate name. Default length is 256.
SNC Quality of Protection (QoP)	<p>Specifies the SAP PSE or certificate name. You can select from the following options:</p> <ul style="list-style-type: none"> - 1 - Apply authentication only. - 2 - Apply integrity protection (authentication). - 3 - Apply privacy protection (integrity and authentication). - 8 - Apply the default protection. - 9 - Apply the maximum protection. <p>Default is 3 - <i>Apply privacy protection (integrity and authentication)</i>.</p>
SAP Cryptographic Library Path	<p>The path to the cryptographic library. Specify sapcrypto.dll for Windows or libsapcrypto.so for Linux.</p>

Connection property	Description
Use X509 Certificate	Specifies the quality of protection. Select to use X509 Certificate based SNC connection.
X509 Certificate Path or SAP Username	<p>The path to the X509 certificate file.</p> <p>If you select to use the X509 certificate, specify the path to the X509 certificate file with .crt extension. You do not need to specify the SAP user name and password.</p> <p>If you do not want to use the X509 certificate, specify the SAP user name for which SNC is configured in SAP Server.</p>
Subscriber Name	<p>A name which defines the Informatica Secure Agent as a unique subscriber in the SAP system.</p> <p>SAP uses this name to define unique operational delta queue (ODQ) when the Secure Agent reads delta data from ODP.</p>
Additional Parameters	<p>Additional SAP parameters that the Secure Agent uses to connect to the SAP system.</p> <p>For example, to generate SAP JCo and SAP CPIC trace, specify the following properties:</p> <pre>jco.client.trace="1"; jco.client.cpic_trace="3";</pre> <p>During the runtime, the JCo and CPIC traces file are generated in the following location:</p> <pre><Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>\ICS\main\bin\rdtm</pre> <p>During the design time, the CPIC traces are generated in the tomcat.out files at the following location:</p> <pre><Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>tomcat.out</pre>
Display Delta Fields	<p>Specifies whether the mapping displays the operation modes that caused the changed data on ODP sources.</p> <p>When enabled, the mapping generates the ODQ_CHANGEMODE and ODQ_ENTITYCNTR fields on the Fields tab for ODP sources that are enabled with Operational Delta Queue (ODQ).</p> <p>Default is disabled.</p>

The following table describes the properties that must configure when you select **Load Balancing Server SNC Connection** as the connection type:

Connection property	Description
SAP Client Number	The client number of the SAP Server.
SAP Language	Language code that corresponds to the SAP language.
SAP Message Server	Host name of the SAP Message Server.
SAP System ID	The system ID of the SAP Message Server.
SAP Group	The login group name, for example, PUBLIC.
SNC My Name	<p>Optional. The Informatica client PSE or certificate name generated on the Secure Agent machine.</p> <p>Default length is 256.</p>

Connection property	Description
SNC Partner Name	The Informatica client PSE or certificate name generated on the SAP Server. Default length is 256.
SNC Quality of Protection (QoP)	Specifies the SAP PSE or certificate name. You can select from the following options: <ul style="list-style-type: none"> - 1 - Apply authentication only. - 2 - Apply integrity protection (authentication). - 3 - Apply privacy protection (integrity and authentication). - 8 - Apply the default protection. - 9 - Apply the maximum protection. Default is 3 - <i>Apply privacy protection (integrity and authentication)</i> .
SAP Cryptographic Library Path	The path to the cryptographic library. Specify <code>sapcrypto.dll</code> for Windows or <code>libsapcrypto.so</code> for Linux.
Use X509 Certificate	Specifies the quality of protection. Select to use X509 Certificate based SNC connection.
X509 Certificate Path or SAP Username	The path to the X509 certificate file. If you select to use the X509 certificate, specify the path to the X509 certificate file with <code>.cert</code> extension. You do not need to specify the SAP user name and password. If you do not want to use the X509 certificate, specify the SAP user name for which SNC is configured in SAP Server.
Subscriber Name	A name which defines the Informatica Secure Agent as a unique subscriber in the SAP system. SAP uses this name to define unique operational delta queue (ODQ) when the Secure Agent reads delta data from ODP.
Additional Parameters	Additional SAP parameters that the Secure Agent uses to connect to the SAP system. For example, to generate SAP JCo and SAP CPIC trace, specify the following properties: <pre>jco.client.trace="1"; jco.client.cpic_trace="3";</pre> During the runtime, the JCo and CPIC traces file are generated in the following location: <pre><Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>\ICS\main\bin\rdtm</pre> During the design time, the CPIC traces are generated in the <code>tomcat.out</code> files at the following location: <pre><Informatica Secure Agent installation directory>\apps \Data_Integration_Server\<DIS version>tomcat.out</pre>
Display Delta Fields	Specifies whether the mapping displays the operation modes that caused the changed data on ODP sources. When enabled, the mapping generates the <code>ODQ_CHANGEMODE</code> and <code>ODQ_ENTITYCNTR</code> fields on the Fields tab for ODP sources that are enabled with Operational Delta Queue (ODQ). Default is disabled.

ServiceNow Mass Ingestion connection properties

When you set up a ServiceNow Mass Ingestion connection, you must configure the connection properties.

The properties of a ServiceNow Mass Ingestion connection vary based on the authentication method that you specify for the connection. When you create a connection, you can select one of the following authentication methods:

- **OAuth 2.0:** Authenticates the connection by using the details of the OAuth API endpoint that is created for the connection in ServiceNow. To use this method, you must create OAuth API endpoint in ServiceNow and then specify the client ID and client secret of the API endpoint in the connection properties. For more information about creating an OAuth API endpoint in ServiceNow, see the [ServiceNow documentation](#).
- **Basic:** Authenticates the connection by validating the login credentials of the ServiceNow account.

Connection properties for OAuth 2.0 authentication

The following table describes the connection properties for a ServiceNow Mass Ingestion connection configured with OAuth 2.0 authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be ServiceNow Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
User Name	User name of the ServiceNow account.
Password	Password for the ServiceNow account.
Client Secret	Client secret of the API endpoint created for the connection in ServiceNow.
Client ID	Client ID of the API endpoint created for the connection in ServiceNow.
Base URI	URL of the ServiceNow instance. You must enter the base URI in the following format: <code>https://{your_servicenow_instance}.service-now.com/</code>
OAuth Token URL	OAuth token endpoint of the ServiceNow instance. The API client associated with the connection sends the access token requests to this endpoint.

Connection properties for Basic authentication

The following table describes the connection properties for a ServiceNow Mass Ingestion connection configured with Basic authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be ServiceNow Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
User Name	User name of the ServiceNow account.
Password	Password for the ServiceNow account.
Base URI	URL of the ServiceNow instance. You must enter the base URI in the following format: <code>https://{your_servicenow_instance}.service-now.com/</code>

Snowflake Data Cloud connection properties

When you set up a Snowflake Data Cloud connection, configure the connection properties.

You can use the following authentication methods to connect to Snowflake:

- Standard. Uses Snowflake account user name and password credentials to connect to Snowflake.
Note: For application ingestion and replication tasks, you can use only the Standard authentication method.
- Authorization Code. Uses the OAuth 2.0 protocol with Authorization Code grant type to connect to Snowflake. Authorization Code allows authorized access to Snowflake without sharing or storing your login credentials.
- KeyPair. Uses the private key file and private key file password, along with the existing Snowflake account user name to connect to Snowflake.
- Client Credentials. Uses the OAuth 2.0 protocol with the Client Credentials grant type to connect to Snowflake. Doesn't apply to application ingestion and replication and database ingestion and replication tasks.

You create a Snowflake Data Cloud connection on the Connections page. You can then use the connection when you read from or write data to Snowflake.

Standard authentication

When you set up a Snowflake Data Cloud connection, configure the connection properties.

The following table describes the Snowflake Data Cloud connection properties for the Standard authentication mode:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + , Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Snowflake Data Cloud connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run application ingestion and replication tasks and database ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Authentication	The authentication method that the connector must use to log in to Snowflake. Select Standard . Default is Standard .
Username	The user name to connect to the Snowflake account.
Password	The password to connect to the Snowflake account.
Account	The name of the Snowflake account. For example, if the Snowflake URL is <code>https://<123abc>.us-east-2.aws.snowflakecomputing.com/console/login#</code> , your account name is the first segment in the URL before <code>snowflakecomputing.com</code> . Here, <code>123abc.us-east-2.aws</code> is your account name. If you use the Snowsight URL, for example, <code>https://app.snowflake.com/us-east-2.aws/<123abc>/dashboard</code> , your account name is <code>123abc.us-east-2.aws</code> Note: Ensure that the account name doesn't contain underscores. To use an alias name, contact Snowflake Customer Support.
Warehouse	The Snowflake warehouse name.
Role	The Snowflake role assigned to the user.
Additional JDBC URL Parameters	The additional JDBC connection parameters. Enter one or more JDBC connection parameters in the following format: <code><param1>=<value>&<param2>=<value>&<param3>=<value>...</code> For example, pass the database and schema values when you connect to Snowflake: <code>db=mydb&schema=public</code> Important: Ensure that there is no space before and after the equal sign (=) when you add the parameters.

OAuth 2.0 authorization code authentication

The following table describes the Snowflake Data Cloud connection properties for an OAuth 2.0 - AuthorizationCode type connection:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + , Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Snowflake Data Cloud connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run application ingestion and replication tasks and database ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Authentication	The authentication method that Snowflake Data Cloud Connector must use to log in to Snowflake. Select AuthorizationCode .
Account	The name of the Snowflake account. For example, if the Snowflake URL is https://<123abc>.us-east-2.aws.snowflakecomputing.com/console/login#, your account name is the first segment in the URL before snowflakecomputing.com. Here, 123abc.us-east-2.aws is your account name. If you use the Snowsight URL, for example, https://app.snowflake.com/us-east-2.aws/<123abc>/dashboard, your account name is 123abc.us-east-2.aws Note: Ensure that the account name doesn't contain underscores. To use an alias name, contact Snowflake Customer Support.
Warehouse	The Snowflake warehouse name.
Additional JDBC URL Parameters	The additional JDBC connection parameters. Enter one or more JDBC connection parameters in the following format: <code><param1>=<value>&<param2>=<value>&<param3>=<value>...</code> For example, pass the database and schema values when you connect to Snowflake: <code>db=mydb&schema=public</code> Important: Ensure that there is no space before and after the equal sign (=) when you add the parameters.
Authorization URL	The Snowflake server endpoint that is used to authorize the user request. The authorization URL is https://<account name>.snowflakecomputing.com/oauth/authorize, where <account name> specifies the full name of your account provided by Snowflake. For example, https://<abc>.snowflakecomputing.com/oauth/authorize Note: If the account name contains underscores, use the alias name. You can also use the Authorization Code grant type that supports the authorization server in a Virtual Private Cloud network.

Property	Description
Access Token URL	<p>The Snowflake access token endpoint that is used to exchange the authorization code for an access token.</p> <p>The access token URL is <code>https://<account name>.snowflakecomputing.com/oauth/token-request</code>, where <account name> specifies the full name of your account provided by Snowflake.</p> <p>For example, <code>https://<abc>.snowflakecomputing.com/oauth/token-request</code></p> <p>Note: If the account name contains underscores, use the alias name.</p>
Client ID	<p>Client ID of your application generated when you create a security integration of type OAuth in Snowflake.</p> <p>For more information, see the Snowflake documentation.</p> <p>Not used by application ingestion and replication and database ingestion and replication tasks.</p>
Client Secret	<p>Client secret generated for the client ID.</p> <p>Not used by application ingestion and replication and database ingestion and replication tasks.</p>
Scope	<p>Determines the access control if the API endpoint has defined custom scopes.</p> <p>Enter space-separated scope attributes.</p> <p>For example, specify <code>session:role:CQA_GCP</code> as the scope to override the value of the default user role. The value must be one of the roles assigned in Security Integration.</p> <p>Not used by application ingestion and replication and database ingestion and replication tasks.</p>
Access Token Parameters	<p>Additional parameters to use with the access token URL.</p> <p>Define the parameters in the JSON format.</p> <p>For example, define the following parameters:</p> <pre>[{"Name": "code_verifier", "Value": "5PMddu6Zcg6Tc4sbg"}]</pre> <p>Not used by application ingestion and replication and database ingestion and replication tasks.</p>
Authorization Code Parameters	<p>Additional parameters to use with the authorization token URL.</p> <p>Define the parameters in the JSON format.</p> <p>For example, define the following parameters:</p> <pre>[{"Name": "code_challenge", "Value": "Ikr-vv52th0UeVRi4"}, {"Name": "code_challenge_method", "Value": "S256"}]</pre> <p>Not used by application ingestion and replication tasks and database ingestion and replication tasks.</p>
Access Token	<p>The access token value.</p> <p>Enter the populated access token value, or click Generate Token to populate the access token value.</p>
Generate Token	<p>Generates the access token and refresh token based on the OAuth attributes you specified.</p>
Refresh Token	<p>The refresh token value.</p> <p>Enter the populated refresh token value, or click Generate Token to populate the refresh token value. If the access token is not valid or expires, the agent fetches a new access token with the help of the refresh token.</p> <p>Note: If the refresh token expires, provide a valid refresh token or regenerate a new refresh token by clicking Generate Token.</p> <p>Not used by application ingestion and replication and database ingestion and replication tasks.</p>

Key pair authentication

The following table describes the Snowflake Data Cloud connection properties for the KeyPair authentication type connection:

Connection property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	The Snowflake Data Cloud connection type.
Runtime Environment	The name of the runtime environment where you want to run the tasks. You cannot run application ingestion and replication tasks and database ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Authentication	The authentication method to log in to Snowflake. Select KeyPair .
Username	The user name to connect to the Snowflake account.
Account	The name of the Snowflake account. For example, if the Snowflake URL is <code>https://<123abc>.us-east-2.aws.snowflakecomputing.com/console/login#/,</code> your account name is the first segment in the URL before <code>snowflakecomputing.com</code> . Here, <code>123abc.us-east-2.aws</code> is your account name. If you use the Snowsight URL, for example, <code>https://app.snowflake.com/us-east-2.aws/<123abc>/dashboard,</code> your account name is <code>123abc.us-east-2.aws</code> . Note: Ensure that the account name doesn't contain underscores. To use an alias name, contact Snowflake Customer Support.
Warehouse	The Snowflake warehouse name.
Additional JDBC URL Parameters	The additional JDBC connection parameters. Enter one or more JDBC connection parameters in the following format: <code><param1>=<value>&<param2>=<value>&<param3>=<value>...</code> For example, pass the database and schema values when you connect to Snowflake: <code>db=mydb&schema=public</code> Important: Ensure that there is no space before and after the equal sign (=) when you add the parameters.
Private Key File	Path to the private key file, including the private key file name, that the Secure Agent uses to access Snowflake. For example, specify the following path and key file name: - On Windows: <code>C:\Users\path_to_key_file\rsa_key.p8</code> - On Linux: <code>/export/home/user/path_to_key_file/rsa_key.p8</code> Note: Verify that the keystore is FIPS-certified.
Private Key Password	Password for the private key file.

Set JDBC URL Parameters

You can configure JDBC URL parameters in the **Additional JDBC URL Parameters** field in the Snowflake Data Cloud connection properties, regardless of which authentication type you select.

Data Ingestion and Replication tasks use only the following parameters:

- Required. To view only the specified database while importing a Snowflake table, enter the database name in the following format:

```
db=<database_name>
```

- Optional. To ignore double-quotation marks in object identifiers and treat all tables as case-insensitive, enter the following parameter:

```
QUOTED_IDENTIFIERS_IGNORE_CASE=true
```

When this property is set to true, Snowflake ignores double-quotation marks in object identifiers and treats all tables as case-insensitive.

Private links to access Snowflake

You can access Snowflake using Azure Private Link endpoints. When you create a Snowflake Data Cloud connection, specify the Snowflake private link account name.

The Azure Private Link setup ensures that the connection to Snowflake uses the Azure internal network and does not take place over the public Internet.

To connect to the Snowflake account over the private Azure network, see [Azure Private Link and Snowflake](#).

Teradata connection properties

When you set up a Teradata connection, you must configure the connection properties.

The following table describes the Teradata connection properties:

Connection property	Description
Connection Name	Name of the connection.
Description	Description of the connection.
Type	Teradata
Runtime Environment	The name of the run-time environment where you want to run the tasks. You cannot use the Hosted Agent for Teradata Connector.
TDPID	The name or IP address of the Teradata database machine.
Tenacity	Amount of time, in hours, that Teradata PT API continues trying to log on when the maximum number of operations runs on the Teradata database. Specify a positive integer. Default is 4.
Database Name	The Teradata database name. If you do not enter a database name, Teradata PT API uses the default login database name.

Connection property	Description
Code Page	<p>Code page associated with the Teradata database.</p> <p>Select one of the following code pages:</p> <ul style="list-style-type: none"> - MS Windows Latin 1. Select for ISO 8859-1 Western European data. - UTF-8. Select for Unicode and non-Unicode data. <p>When you run a task that extracts data from a Teradata source, the code page of the Teradata PT API connection must be the same as the code page of the Teradata source.</p>
Max Sessions	<p>Maximum number of sessions that Teradata PT API establishes with the Teradata database. Specify a positive, non-zero integer. Default is 4.</p>
Min Sessions	<p>Minimum number of Teradata PT API sessions required for the Teradata PT API job to continue.</p> <p>Specify a positive integer between 1 and the Max Sessions value.</p> <p>Default is 1.</p>
Sleep	<p>Amount of time, in minutes, that Teradata PT API pauses before it retries to log on when the maximum number of operations run on the Teradata database.</p> <p>Specify a positive, non-zero integer. Default is 6.</p>
Data Encryption	<p>Enables full security encryption of SQL requests, responses, and data.</p> <p>Default is disabled.</p>
Block Size	<p>Maximum block size, in bytes.</p> <p>Teradata PT API uses this property to read the data block size from source through the Export operator.</p> <p>Maximum is 16775168 bytes for Teradata Database version 16.20 and higher.</p> <p>If the Teradata Database version is lower than 16.20, then Teradata scales down the block size from 16775168 bytes to the maximum allowed value. The block size 16775168 is not allowed in the Spool mode. For more information, see Teradata logs and verify the Teradata documentation of the same version.</p>
Authentication Type	<p>Method to authenticate the user.</p> <p>Select one of the following authentication types:</p> <ul style="list-style-type: none"> - Native. Authenticates your user name and password against the Teradata database specified in the connection. - LDAP. Authenticates user credentials against the external LDAP directory service. - KRB5. Authenticates to the Teradata database through Kerberos. <p>Default is Native.</p>
Kerberos Artifacts Directory	<p>Directory that contains Kerberos configuration files named <code>krb5.conf</code> and <code>IICSTPT.keytab</code>. Applicable when you select KRB5 as the authentication type.</p>
Metadata Advanced Connection Properties	<p>The values to set the optional properties of the JDBC driver to fetch the metadata.</p> <p>For example, <code>tmode=ANSI</code>.</p>
Enable Metadata Quotification	<p>Select this option to enable the Teradata connection to read reserved words used as table or column names from the Teradata database.</p> <p>By default, the Enable Metadata Quotification checkbox is not selected and the Secure Agent does not read reserved words from Teradata.</p>

Connection property	Description
User Name	Database user name with the appropriate read and write database permissions to access the database. If you select KRB5 as the authentication type, you must specify the Kerberos user name.
Password	Password for the database user name. If you select KRB5 as the authentication type, you do not need to specify the Kerberos user password.

Workday Mass Ingestion connection properties

When you set up a Workday Mass Ingestion connection, you must configure the connection properties.

The properties of a Workday Mass Ingestion connection vary based on the authentication method that you specify for the connection. When you create a connection, you can select one of the following authentication methods:

- **Basic:** Authenticates the connection by validating the login credentials of the Workday account.
- **OAuth 2.0 Refresh Token Flow:** Authenticates the connection by using an application that is registered in Workday. To use this method, you must register an application in Workday and then specify the client ID and client secret of the application in the connection properties. For more information about registering an application in Workday, see the [Workday documentation](#).

Connection properties for Basic authentication

The following table describes the connection properties for a Workday Mass Ingestion connection configured with Basic authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Workday Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Domain Name	Name of the Workday domain that contains the resources that you want to access.
Tenant Name	Identifier of the Workday tenant that you want to access.

Connection property	Description
Version	Optional. Web Service Description Language (WSDL) version for the endpoints that the connection must use to retrieve Workday data. The list of operations supported for a web service depends on the WSDL version that you specify in this field. Note: Informatica recommends that you use WSDL v37.0 because Workday Mass Ingestion connections might not read data from the services that are not part of WSDL v37.0. For more information on the WSDL versions, see the Workday Web Services (WWS) documentation .
User Name	User name of the Workday account.
Password	Password for the Workday account.

Note: If you configure a connection with the Basic authentication method and then test the connection, the test is always successful even if the connection property values that you specified are incorrect. Therefore, ensure that you specify correct values for the connection properties before you save the connection.

Connection properties for OAuth 2.0 Refresh Token Flow authentication

The following table describes the connection properties for a Workday Mass Ingestion connection configured with OAuth 2.0 Refresh Token Flow authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Workday Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Domain Name	Name of the Workday domain that contains the resources that you want to access.
Tenant Name	Identifier of the Workday tenant that you want to access.
Version	Optional. Web Service Description Language (WSDL) version for the endpoints that the connection must use to retrieve Workday data. The list of operations supported for a web service depends on the WSDL version that you specify in this field. Note: Informatica recommends that you use WSDL v37.0 because Workday Mass Ingestion connections might not read data from the services that are not part of WSDL v37.0. For more information on the WSDL versions, see the Workday Web Services (WWS) documentation .
User Name	Optional. User name of the Workday account.
Client ID	Client ID of the application registered in Workday.

Connection property	Description
Client Secret	Private key of the application registered in Workday.
Refresh Token	Refresh token string that Workday generates for the registered application.
Token Endpoint	OAuth token endpoint of the Workday instance. The registered application sends the access token requests to this endpoint.

Zendesk Mass Ingestion connection properties

When you set up a Zendesk Mass Ingestion connection, you must configure the connection properties.

The properties of a Zendesk Mass Ingestion connection vary based on the authentication method that you specify for the connection. When you create a connection, you can select one of the following authentication methods:

- **Basic:** Authenticates the connection by using the login credentials and subdomain associated with the Zendesk account. The Basic authentication method does not use any encrypted access token to connect to the data source, which results in quick and easy access to Zendesk data.

Note: You can use the Basic authentication method only if your Zendesk account is not configured with two-factor authentication. If the account is configured with two-factor authentication, you must use the OAuth 2.0 authentication method for the connection.

- **OAuth 2.0:** Authenticates the connection by using an application that is registered in Zendesk along with the login credentials and subdomain associated with the Zendesk account. To use this method, you must register an application in Zendesk and then specify the client ID and client secret of the application in the connection properties. For more information about registering an application in Zendesk, see the [Zendesk documentation](#).

Connection properties for Basic authentication

The following table describes the connection properties for a Zendesk Mass Ingestion connection configured with Basic authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Zendesk Mass Ingestion .

Connection property	Description
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Email ID	User name of the Zendesk account. The user name is an email address.
Password	Password for the Zendesk account.
Subdomain	URL of the Zendesk help center that you want to access.

Note: For more information about the Basic authentication method, see the Zendesk documentation.

Connection properties for OAuth 2.0 authentication

The following table describes the connection properties for a Zendesk Mass Ingestion connection configured with OAuth 2.0 authentication:

Connection property	Description
Connection Name	A name for the connection. This name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + - Spaces at the beginning or end of the name are trimmed and are not saved as part of the name. Maximum length is 100 characters. Connection names are not case sensitive.
Description	An optional description for the connection. Maximum length is 255 characters.
Type	The type of connection. For an Oracle Database Ingestion connection, the type must be Zendesk Mass Ingestion .
Runtime Environment	Name of the runtime environment where you want to run the ingestion tasks. You must specify a Secure Agent as the runtime environment. Note: You cannot run application ingestion and replication tasks on a Hosted Agent or serverless runtime environment.
Email ID	User name of the Zendesk account. The user name is an email address.
Password	Password for the Zendesk account.
Subdomain	URL of the Zendesk help center that you want the connection to access.
Client ID	Client ID of the application registered in Zendesk.
Client Secret	Client secret of the application registered in Zendesk.
Grant Type	OAuth 2.0 grant type to be used by the connection. By default, Zendesk Mass Ingestion connections are configured to use the password grant type to exchange user names and passwords for access tokens.

Note: For more information about the OAuth 2.0 authentication method, see the Zendesk documentation.

INDEX

A

- Adobe Analytics Mass Ingestion connections
 - connection properties [18](#)
- Advanced FTP V2 connections
 - properties [19](#)
- Advanced FTPS V2 connections
 - properties [21](#)
- Advanced SFTP V2 connections
 - properties [23](#)
- Amazon Kinesis
 - AWS credential profile [28](#)
- Amazon Kinesis connection
 - overview [25](#)
- Amazon Redshift V2
 - connection properties [38](#)
- Amazon Redshift V2 connections
 - overview [28](#)
- Amazon S3 V2
 - connection properties [49](#)
- application ingestion and replication tasks
 - connectors [9](#)
- authentication
 - OAuth 2.0 authorization code [157](#), [159](#)
- Azure Data Lake Storage Gen2
 - connection properties [94](#)

C

- Cloud Application Integration community
 - URL [6](#)
- Cloud Developer community
 - URL [6](#)
- Cloud Integration Hub connections
 - connection properties [56](#)
- connection
 - Amazon Kinesis Firehose
 - connection properties [25](#)
 - Amazon Kinesis Streams
 - connection properties [27](#)
 - properties [56](#)
- connections
 - Adobe Analytics Mass Ingestion [18](#)
 - Amazon Redshift V2 [38](#)
 - Amazon S3 V2 [49](#)
 - AMQP
 - connection properties [54](#)
 - Azure Data Lake Storage Gen2 [94](#)
 - Azure Event Hub
 - connection properties [96](#)
 - Cloud Integration Hub [56](#)
 - configuring for Data Ingestion and Replication tasks [17](#)
 - Data Ingestion and Replication connection properties [17](#)
 - Db2 for i Database Ingestion [67](#)
 - Db2 for LUW Database Ingestion connection [68](#)

- connections (*continued*)
 - Db2 for zOS Database Ingestion [69](#)
 - flat file [70](#)
 - Google Analytics Mass Ingestion [73](#)
 - Google BigQuery [73](#)
 - Google Cloud Storage V2 [75](#)
 - Google PubSub [79](#)
 - JDBC V2 [81](#)
 - JMS
 - connection properties [83](#)
 - Kafka
 - connection properties [84](#)
 - Marketo V3 [92](#)
 - Microsoft Azure Blob Storage V3 [93](#)
 - Microsoft Azure Synapse Analytics Database Ingestion [97](#)
 - Microsoft Azure Synapse SQL [98](#)
 - Microsoft Dynamics 365 Mass Ingestion [100](#)
 - Microsoft SQL Server [104](#)
 - MongoDB Mass Ingestion [107](#)
 - MQTT
 - connection properties [108](#)
 - MySQL [110](#)
 - Netezza [111](#)
 - NetSuite Mass Ingestion [112](#)
 - OPC UA [114](#)
 - Oracle Database Ingestion [117](#)
 - Oracle Fusion Cloud Mass Ingestion [123](#)
 - PostgreSQL [124](#)
 - REST V2 [126](#)
 - Salesforce Marketing Cloud [136](#)
 - Salesforce Mass Ingestion [138](#)
 - SAP HANA Database Ingestion [140](#)
 - SAP Mass Ingestion [143](#)
 - SAP ODP Extractor [149](#)
 - ServiceNow Mass Ingestion [154](#)
 - Snowflake Data Cloud [155](#), [156](#)
 - Teradata connection [160](#)
 - testing for Data Ingestion and Replication tasks [17](#)
 - Workday Mass Ingestion [162](#)
 - Zendesk Mass Ingestion [164](#)
- connections Hadoop Files V2 [80](#)
- connectors
 - application ingestion and replication tasks [9](#)
 - database ingestion and replication tasks [11](#)
 - for Data Ingestion and Replication tasks [9](#)

D

- Data Ingestion and Replication connections
 - connection properties [17](#)
- Data Ingestion and Replication connectors
 - overview [9](#)
- Data Integration community
 - URL [6](#)

database ingestion and replication tasks
connectors [11](#)
Db2 for i Database Ingestion connections
connection properties [67](#)
Db2 for LUW Database Ingestion connection
connection properties [68](#)
Db2 for zOS Database Ingestion connections
connection properties [69](#)

F

flat file
connection properties [70](#)

G

Google Analytics Mass Ingestion connections
connection properties [73](#)
Google BigQuery
connection properties [73](#)
Google Cloud Storage V2
connection properties [75](#)
Google PubSub
connection properties [79](#)

H

Hadoop Files V2
connection properties [80](#)

I

Informatica Global Customer Support
contact information [7](#)
Informatica Intelligent Cloud Services
web site [6](#)

J

JDBC V2
connection properties [81](#)

K

Kerберised Kafka
prerequisites [87](#)

L

Linux
configuring proxy settings [77](#)

M

maintenance outages [7](#)
Marketo V3
connection properties [92](#)
Microsoft Azure Blob Storage V3
connection properties [93](#)

Microsoft Azure Synapse Analytics Database Ingestion connections
connection properties [97](#)
Microsoft Azure Synapse SQL
connection properties [98](#)
Microsoft Dynamics 365 Mass Ingestion connections
connection properties [100](#)
Microsoft Fabric OneLake
connection properties [106](#)
Microsoft SQL Server
connection properties [104](#)
MongoDB Mass Ingestion
connection properties [107](#)
MySQL
connection properties [110](#)

N

Netezza
connection properties [111](#)
NetSuite Mass Ingestion connections
connection properties [112](#)

O

OPC UA
connection properties [114](#)
Oracle Cloud Object Storage connections
properties [116](#)
Oracle Database Ingestion connections
connection properties [117](#)
Oracle Fusion Cloud Mass Ingestion connections
connection properties [123](#)

P

PostgreSQL
connection properties [124](#)
proxy settings
configuring on Linux [77](#)
configuring on Windows [76](#)

R

REST V2
authentication
standard [126](#)
connection properties [126](#)

S

Salesforce Marketing Cloud
connection properties [136](#)
Salesforce Mass Ingestion connections
connection properties [138](#)
SAP HANA Database Ingestion connections
connection properties [140](#)
SAP Mass Ingestion connections
connection properties [143](#)
ServiceNow Mass Ingestion connections
connection properties [154](#)
Snowflake Data Cloud
authentication
standard [156](#)

Snowflake Data Cloud (*continued*)
 connection properties [155](#), [156](#)
status
 Informatica Intelligent Cloud Services [7](#)
system status [7](#)

T

Teradata connection
 connection properties [160](#)
trust site
 description [7](#)

U

upgrade notifications [7](#)

W

web site [6](#)
Windows
 configuring proxy settings [76](#)
Workday Mass Ingestion connections
 connection properties [162](#)

Z

Zendesk Mass Ingestion connections
 connection properties [164](#)