Informatica® Cloud Data Integration
July 2024

VSAM Connector

Informatica Cloud Data Integration VSAM Connector
July 2024

© Copyright Informatica LLC 2022, 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, the Informatica logo, Informatica Cloud, and PowerCenter 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.

See patents at https://www.informatica.com/legal/patents.html.

DISCLAIMER: Informatica LLC provides this documentation "as is" without warranty of any kind, either express or implied, including, but not limited to, the implied warranties of noninfringement, merchantability, or use for a particular purpose. Informatica LLC does not warrant that this software or documentation is error free. The information provided in this software or documentation may include technical inaccuracies or typographical errors. The information in this software and documentation is subject to change at any time without notice.

NOTICES

This Informatica product (the "Software") includes certain drivers (the "DataDirect Drivers") from DataDirect Technologies, an operating company of Progress Software Corporation ("DataDirect") which are subject to the following terms and conditions:

1. THE DATADIRECT DRIVERS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.
2. IN NO EVENT WILL DATADIRECT OR ITS THIRD PARTY SUPPLIERS BE LIABLE TO THE END-USER CUSTOMER FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR OTHER DAMAGES ARISING OUT OF THE USE OF THE ODBC DRIVERS, WHETHER OR NOT INFORMED OF THE POSSIBILITIES OF DAMAGES IN ADVANCE. THESE LIMITATIONS APPLY TO ALL CAUSES OF ACTION, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS.

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 5

Informatica Resources 5

- Informatica Documentation 5
- Informatica Intelligent Cloud Services web site 5
- Informatica Intelligent Cloud Services Communities 5
- Informatica Intelligent Cloud Services Marketplace 6
- Data Integration connector documentation 6
- Informatica Knowledge Base 6
- Informatica Intelligent Cloud Services Trust Center 6
- Informatica Global Customer Support 6

Chapter 1: Introduction to VSAM Connector 7

VSAM Connector usage scenario 7

Integration with the PowerExchange environment 9

VSAM Connector asset types 10

Software requirements 11

Supported targets 11

VSAM Connector rules and guidelines 12

Prerequisite tasks for using VSAM Connector 12

Task flow for using VSAM Connector 13

Chapter 2: VSAM connections 14

VSAM connection properties 14

Chapter 3: Mappings and mapping tasks with VSAM Connector 17

Mapping configuration 17

- VSAM source properties in mappings 18
- Targets in mappings with VSAM sources 22
- VSAM targets in a mapping 22
- Lookup Transformation 24
- Multiple-record writes to nonrelational targets 25

Mapping task configuration 25

- Properties for parameterized sources and connections 25

Running a mapping task 26

Monitoring a mapping task 26

Chapter 4: VSAM administration 28

Upgrading to a new version of PowerExchange 28

Chapter 5: Data type reference **29**

VSAM and transformation data types 29

Chapter 6: Connection overrides reference **32**

Connection overrides 32

Index **33**

Preface

Learn how to use the *VSAM Connector* to retrieve VSAM records from a PowerExchange environment and write the data to a supported target. The VSAM connector works with Cloud Data Integration mappings, mapping tasks, and taskflows. Learn how to configure connections, mappings, and mapping tasks and to run and monitor mapping tasks.

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

Introduction to VSAM Connector

You can use VSAM Connector to connect to a PowerExchange VSAM environment to retrieve records for VSAM source data sets and write the data in batch to a supported target. The PowerExchange Listener retrieves metadata from the data map repository and records from the VSAM data sets. The data is returned to the PowerExchange Bulk Reader, which sends the records to Data Integration. Data Integration can then send the records to a supported target.

VSAM Connector must be able to connect to an active PowerExchange environment that includes the PowerExchange Listener. You do not need to copy the PowerExchange dbmover configuration file to the VSAM Connector system. For more information about configuring PowerExchange for VSAM Connector, see the *PowerExchange Bulk Data Movement Guide*.

VSAM Connector works with the Informatica Cloud Secure Agent. When you install and start the Secure Agent, it downloads the VSAM Connector components and the PowerExchangeClient packages. When you define a runtime environment, you configure a Secure Agent group.

To retrieve source records from the data map repository, you run a mapping or a mapping task.

When you define mappings in Data Integration, you select the data maps that PowerExchange generated for the VSAM source data sets. The VSAM Connector metadata adapter can then retrieve the source metadata from the data maps and store this information in the cloud.

After you materialize a target, you can use PowerExchange Change Data Capture (CDC), if this option is available at your site, to update the target on a continuous basis. By moving only the change data, you can refresh data on the target faster.

Note: VSAM Connector does not work with the Hosted Agent.

VSAM Connector usage scenario

VSAM Connector uses group source processing to read bulk data from ESDS, KSDS, or RRDS data sets with multiple record types.

Group source processing reads data that is stored in the same physical source in a single pass. By eliminating multiple passes of the source data, group source processing enhances bulk data movement performance and reduces resource consumption on the source system.

You can use the browser-based Informatica Intelligent Cloud Services interfaces to define VSAM connections, mappings, and mapping tasks.

With Informatica Intelligent Cloud Services interfaces, you gain access to easy-to-use, browser-based interfaces for defining mappings and mapping tasks and running and monitoring mapping tasks. When you run the mapping and mapping tasks, VSAM Connector components, in conjunction with Data Integration, read and process the records that VSAM Connector read from the VSAM data set.

VSAM Connector can move bulk data from the following types of VSAM data sets:

Entry-sequenced data set (ESDS)

A data set that stores records in the order in which they were created, with the latest records at the end. ESDSs have no index. ESDS records can be accessed directly by using a relative byte address (RBA) or sequentially based on entry sequence. VSAM Connector accesses ESDS data sets sequentially.

Key-sequenced data set (KSDS)

A data set that contains records that are loaded in key sequence and controlled by an index. KSDSs can be accessed directly by using a key or RBA or sequentially based on key sequence.

By default, VSAM Connector accesses KSDS records sequentially if no WHERE clause is specified.

VSAM Connector accesses records by using a key under the following conditions:

- A WHERE clause is specified.
- The primary field in the KSDS key is used in one of more WHERE clause conditions.

For example, VSAM Connector uses keyed access to perform a PowerCenter Lookup transformation that uses a KSDS lookup source when the lookup condition includes question marks that are substituted at run time. VSAM Connector always uses a key to access a KSDS file if possible when processing a WHERE clause.

VSAM Connector can access KSDS files through an alternate index that you specify by using a file name override. VSAM Connector determines whether to use keyed or sequential access at run time after it opens the file and determines the start position and length of the key.

When you define the data map, you can optimize the data map for sequential access by enabling CI access and setting the number of data buffers to a high value. You can optimize the data map for keyed access by setting the number of index buffers to a high value and the Number of data buffers to a low value, such as 2.

Relative record data set (RRDS)

A data set that contains fixed-length records. RRDSs are accessed sequentially or directly by using a relative record number (RRN). VSAM Connector accesses RRDS data sets sequentially based on RRN.

Variable-length relative record data set (VRRDS)

A data set that contains variable-length records. VRRDSs can be accessed sequentially or directly by using a relative record number (RRN). VSAM Connector accesses VRRDS records sequentially based on RRN.

PowerExchange can process VSAM data sets that are in basic or extended format. PowerExchange can also processes data sets that contain compressed data.

Because VSAM is not a relational database, you must create a data map in the PowerExchange Navigator. VSAM Connector uses the data map to access VSAM data and metadata and to create a relational row-type view of the records. VSAM Connector requires a relational view to use SQL-type statements to read or write bulk data.

When you create VSAM data maps, you must select the access method that corresponds to the VSAM data set type. The following VSAM access methods are available:

- ESDS
- KSDS
- RRDS, for both fixed-length and variable-length relative record data sets

If a copybook is available for the VSAM data set, you can import it for the data map to get a detailed layout of each record. If you do not import a copybook, you will need to manually define the fields in the VSAM records.

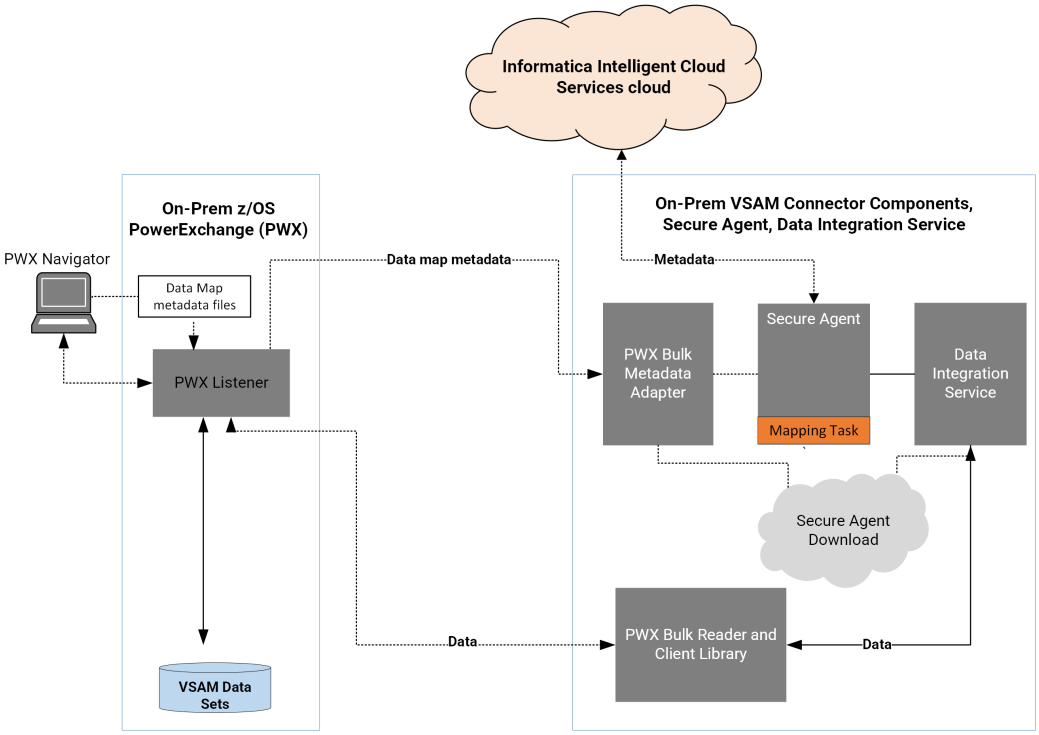
Also, you must run a PowerExchange Listener on the MVS system where the VSAM source or target is located. Verify that this PowerExchange Listener can communicate with the system or systems where the PowerExchange Navigator, PowerCenter Client, and PowerCenter Integration Service run.

Integration with the PowerExchange environment

You must have a functional on-premises PowerExchange environment that can retrieve records from VSAM source data sets on a z/OS system. The VSAM Connector retrieves records from the data map metadata.

You must have a license that allows your organization to use the VSAM Connector and PowerExchangeClient packages on a z/OS system. On each system from which you plan to use VSAM Connector, install the Secure Agent. After you start the Secure Agent the first time, the VSAM Connector and PowerExchangeClient package are installed locally. You then can use Cloud Data Integration, Administrator, and Monitor interfaces to configure VSAM connections, mappings, and mapping tasks and to run and monitor mapping tasks.

The following image shows the general architecture of the VSAM Connector components in relation to Cloud Data Integration and the remote VSAM source system:



This configuration spans the following on-premises systems:

- The Windows system where the PowerExchange Navigator client runs.
- The z/OS system where the VSAM data sets and PowerExchange Listener run.
- The system where the VSAM Connector components and Secure Agent run.

Based on this configuration, the following processing occurs:

1. PowerExchange retrieves records from VSAM data sets and sends the data to the PowerExchange Listener for z/OS.

   You defined a source transformation for each VSAM source table in the PowerExchange Navigator. The PowerExchange Navigator generated a corresponding data map, which will be used to retrieve the records.

2. When you design a mapping in Data Integration, you associate the Source transformation with a VSAM connection. You can define the connection in either Data Integration or Administrator. The PWX Bulk Metadata Adapter then uses the connection to connect to the PowerExchange Listener on z/OS and retrieve data map metadata files for the VSAM source from the PowerExchange Listener.

   In the mapping, you must also define a Target transformation and connection and map the source fields to target fields.

3. The Secure Agent in conjunction with Data Integration sends the data map metadata to the Informatica Intelligent Cloud Services repository in the cloud.

4. When you run the mapping task that is associated with the mapping, all of the metadata, connection properties, mapping information, advanced source properties, runtime properties, and advanced session properties are pushed down to the on-premises Secure Agent in the form of a session .xml file.

5. The PWX Bulk Reader on the Secure Agent machine communicates with the PowerExchange Listener on the z/OS machine to retrieve the records from the PowerExchange Listener.

6. Data Integration transmits the records to the mapped target.

Alternative configurations of PowerExchange with the VSAM Connector components and Cloud Data Integration are possible. For example:

- The PowerExchange Navigator can run on any supported Windows system, including a Windows system where the VSAM Connector and Secure Agent run.

VSAM Connector asset types

The following table lists the Data Integration asset types that the VSAM Connector uses:

Asset Type	Required or Optional?
Mapping	Required
Mapping task	Optional
Taskflow	Optional

Software requirements

Verify that your PowerExchange system and VSAM Connector system meet the following software requirements for integration with Informatica Cloud Data Integration:

- You must have a remote PowerExchange installation on the VSAM source system. You also must have a local or remote PowerExchange installation wherever you plan to run the PowerExchange Navigator. For all PowerExchange installations, the following PowerExchange versions are supported:
  - PowerExchange 10.5.*x*
  - PowerExchange 10.4.1
  - PowerExchange 10.4.0 HotFix 1
  - PowerExchange 10.4.0
- VSAM Connector components run on the same Linux and Windows operating systems that are supported for the Informatica Secure Agent. For more information, see the Product Availability Matrix (PAM) for Informatica Intelligent Cloud Services (IICS) <*current_version*>, which is available on Informatica Network at https://network.informatica.com/community/informatica-network/product-availability-matrices. You must have a login to access this location.
- The supported versions of VSAM source databases and operating systems are the same as those that PowerExchange supports. For more information, see the **PWX Classic** tab in the Product Availability Matrix (PAM) for Informatica <current_pwx_version>, which is available on Informatica Network at https://network.informatica.com/community/informatica-network/product-availability-matrices.

Supported targets

VSAM Connector is certified with the following target types:

- Amazon Redshift
- Aurora PostgreSQL
- Db2 for z/OS
- Db2 for i
- Google BigQuery
- Microsoft SQL Server and Amazon Relational Database Service (RDS) for SQL Server
- MySQL
- Oracle and Amazon RDS for Oracle
- PostgreSQL
- Snowflake

Important: Other Cloud Data Integration targets are expected to work but have not been fully tested and certified with this source type. Informatica recommends that you use other targets in your test environment first.

VSAM Connector rules and guidelines

Review the VSAM Connector rules and guidelines before you begin configuring connections, mappings, and mapping tasks.

- You can use VSAM connections for sources and targets.
- VSAM Connector connects to the PowerExchange Listener by using the **Listener Location** property that you specify in the VSAM connection definition. On the Connector system, you do not need to create a local copy of the DBMOVER configuration file that contains a NODE statement pointing to the PowerExchange Listener.
- If you upgrade your system to a later PowerExchange version, you must stop the Secure Agent that connects to the older PowerExchange Listener before you try to connect to the upgraded PowerExchange Listener. After you edit the connection properties to specify the upgraded PowerExchange Listener, restart the Secure Agent before you try to establish the connection again. For more information, see "Upgrading to a new version of PowerExchange" on page 28.
- The PowerExchange Listener on the LUW system can compress or encrypt the data that it sends to the PWX Bulk Reader component of VSAM Connector depending on how you set the **Compression** and **Encryption** properties for the source connection.
- VSAM Connector uses group source processing to read bulk data from ESDS, KSDS, or RRDS data sets with multiple record types. Group source processing reads data that is stored in the same physical source in a single pass. By eliminating multiple passes of the source data, group source processing enhances bulk data movement performance and reduces resource consumption on the source system.

Prerequisite tasks for using VSAM Connector

Before you begin configuring VSAM connections, mappings, and mapping tasks, you must complete several prerequisite tasks in PowerExchange and in Informatica Intelligent Cloud Services Administrator. Make sure you have a user role that allows you to perform these tasks.

- Verify that the PowerExchange installations on the z/OS source system and PowerExchange Navigator system use a PowerExchange version that the VSAM Connector supports. Test your configuration by verifying that PowerExchange is able to retrieve records from VSAM source data sets.
- If your site uses TLS or SSL security, verify that valid certificates are available for the source and target servers and that the PWX_CONFIG environment variable points to the dbmover configuration file that includes the SSL configuration for the PowerExchange Listener. For more information about configuring TLS or SSL security, see the *PowerExchange Reference Manual*.
- If you localize PowerExchange messages in Japanese, before you start the Secure Agent, create a copy of the PowerExchange dbmover_ja.cfg file. Then set the PWX_CONFIG environment variable to point to the copy. For more information about the dbmover_ja.cfg configuration file, see the *PowerExchange Reference Manual*.
- Download the Secure Agent. The Secure Agent can then download the VSAM Connector components and PowerExchangeClient packages based on your licensing. The download also creates a Secure Agent group that contains a Secure Agent as the runtime environment. You can add another Secure Agent to the group. You must have at least one runtime environment in each organization so that users in the organization can run tasks.

  Note: VSAM Connector does not support the Hosted Agent.

- In Administrator, configure user roles if the predefined roles are not sufficient. Then assign one or more roles to each user or user group. The roles must provide the privileges that the users require to perform all of their tasks. Usually, the organization administrator configures user roles.
- If you do not want users to run scheduled mapping tasks during a certain period, define a blackout period for all schedules on the **Schedules** page in Administrator.

Task flow for using VSAM Connector

To use VSAM Connector, create a source connection, mapping, and mapping task. Also create a target connection if one is not already defined.

Before you begin, in the Administrator, you must complete the prerequisite tasks. See "Prerequisite tasks for using VSAM Connector" on page 12.

1. In Administrator, create a source connection with the type of **VSAM**. Test the connection.

   Alternatively, you can create a source connection in Data Integration when you configure a mapping. When you create a connection, you must set some connection properties that are specific to VSAM Connector. For more information, see "VSAM connection properties" on page 14.

   VSAM Connector uses this connection to retrieve source metadata and data records from PowerExchange.
2. Create a target connection with a supported target type, if one is not already defined.
3. In Data Integration, configure a mapping. Validate the mapping and then save it.
4. Optional. Create a mapping task that is associated with the mapping.
5. Run the mapping task manually, or create a schedule to run it on a recurring basis. For more information, see "Running a mapping task" on page 26.
6. Monitor the job status for the mapping task instance. For more information, see "Monitoring a mapping task" on page 26.

   If a job has the status of Failed or Warning, check the session log for related error or informational messages. Then try to correct the issues that caused the failure and restart the mapping task from the appropriate point in the stream.

CHAPTER 2

VSAM connections

You must create a VSAM connection to connect to a PowerExchange Listener to read from one or more VSAM source data sets.

The PWX Bulk Reader that is on the system where the Secure Agent runs requires the connection information to retrieve VSAM records from PowerExchange.

You can create a VSAM connection in Administrator or Data Integration, provided that your assigned user role allows this action. Use one of the following methods:

- In Administrator, click **Connections** on the navigation bar. On the **Connections** page, click **New Connection** and then enter connection details and properties for a connection that has the type of **VSAM**.
- In Data Integration, when you create a mapping, click **Source** on the mapping canvas and then click **Source** in the **Properties: Source** panel. To create a connection, click **New Connection** and define a connection that has the type of **VSAM**. Alternatively, if you previously defined a VSAM connection, select it from the **Connection** list.

When you configure a connection, you must specify a runtime environment for the connection. The runtime environment must contain a Secure Agent with the status "Up and Running."

After you create a VSAM connection, it becomes available for use within the organization and can be used in multiple mappings.

VSAM connection properties

When you configure a VSAM connection, you must set the connection properties.

The following table describes the VSAM connection properties:

Property	Description
Connection Name	A name for the VSAM 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	Description of the VSAM connection. Maximum length is 4000 characters.
Type	Type of connection. For VSAM, the type must be **VSAM**.

Property	Description
Runtime Environment	Name of the runtime environment that contains the Secure Agent that you want to use to run mapping tasks.
Listener Location	Host name or IP address of the system where the PowerExchange Listener that processes requests for VSAM runs. Also include the Listener port number. Enter the value in the following format, where *host_name* can be a host name or IP address: *host_name*:*port_number* For example: LSNR1:1467
User Name	A user name that can be used to access the PowerExchange Listener when PowerExchange Listener security is enabled. For more information, see the SECURITY statement in the *PowerExchange Reference Manual*.
Password	Password that is associated with the user name that is specified in the **User Name** property.
Schema Name	The schema name of the data map.
Code Page	The code page that the Secure Agent for the Data Integration Service uses to extract data from the source file.
Offload Processing	Controls whether to use offload processing. Offload processes transfers bulk data processing from the source system to the target system. Options are: - **Auto**. Cloud Data Integration determines whether to use offload processing. - **Filter After**. Offloads the bulk data processing to the target, including the filtering of data. - **Filter Before**. Offloads processing to the target but continues to filter data on the source system. - **No**. Disables offload processing. Default is No.
Offload Threads	The number of threads that Cloud Data Integration uses to process bulk data. For optimal performance, this value should not exceed the number of installed or available processors on the machine where the secure agent runs. Valid values are 1 through 64. Default is 0, which disables multithreading. Not all connection types support offload threads. If the **Offload Threads** connection attribute for one of these connections is set to a nonzero value, processing continues without threads.
Array Size	For VSAM data sets and sequential files, the size of the storage array, in number of records, that is used for partitioned or multithreaded sessions. For partitioned sessions, this array size is shared across the partitions. For multithreaded sessions, each thread has this array size. Valid values are from 1 through 5000. Default is 25. To tune partitioned sessions, particularly when the **Write Mode** attribute specifies **Confirm Write On**, increase the array size.
Replace Low Values with Spaces	Controls whether to replace embedded nulls in character data with spaces. Select this property to replace embedded nulls. By default, this property is selected.
Connection Retry Period	Number of seconds that the PowerExchange Bulk Reader tries to reconnect to the PowerExchange Listener after the initial connection attempt fails. If a connection cannot be established within the retry period, the mapping task fails. The default is 0, which disables connection retries.

Property	Description
Custom Properties	Custom properties or connection property overrides. Custom properties are properties that you can specify to override PowerExchange default settings. You can enter multiple properties by using a semicolon (;) as the separator. For example: `<property>=<value>;<property>=<value>` Normally, you set custom properties only at the direction of Informatica Global Customer Support. **Note:** These properties are equivalent to the **PWX Override** options for the PowerExchange Client for PWX NRDB Batch connections in PowerCenter. You can also specify connection property overrides in this field or through a parameter file. To specify connection property overrides through a parameter file, you set a parameter in the form of `$<ParameterName>`, where you prefix a user-defined parameter name with a dollar sign character ($). Then configure a mapping task for the mapping to use a parameter file that contains the user-defined parameter definition by specifying the parameter file name in the **Parameter File Name** field on the **Runtime Options** tab. **Note:** - If you enter the same parameter for both the mapping and connection, the connection custom property takes precedence. - If you have a parameter file, the parameter name you specify in this field must match an entry defined in the parameter file. For more information, see the “Connection overrides reference” chapter.
Write Properties	Write Mode. Options are: - **Confirm Write On**. Sends data to the PowerExchange Listener and waits for a success or no success response before sending more data. This mode sends data synchronously to the PowerExchange Listener rather than buffering the data. - **Confirm Write Off**. Sends data asynchronously to the PowerExchange Listener by buffering the data. This mode does not wait for a success or no success response.

CHAPTER 3

Mappings and mapping tasks with VSAM Connector

In Data Integration, create a mapping that includes a VSAM source and a target of a supported type. Then you can optionally create a mapping task that is associated with the mapping to process data based on the mapping.

The mapping and mapping task information is stored in the Informatica Intelligent Cloud Services cloud.

To process the VSAM records, run the mapping task in your runtime environment. The mapping task runs on premises, on the same machine where the Secure Agent and PowerExchange client components run. The mapping task uses the mapping and mapping task information pushed down from the Informatica Intelligent Cloud Services cloud. During a mapping task run, the PWX Bulk Reader retrieves the VSAM records and Data Integration transmits the records to the target based on the mapping.

After a mapping task ends, you can determine its job status and view the session log from either the Data Integration or Monitor interface.

Mapping configuration

You must configure a mapping that contains a Source transformation for a VSAM source and map it to a Target transformation.

To begin creating a mapping in Data Integration, click **New** > **Mappings** > **Mapping** and then click **Create**. In the **New Mapping** dialog box, click **Mapping** > **Continue**.

When you define a Source transformation in a mapping, specify a name for the transformation, a VSAM connection, the source type, and the data map tables for the VSAM source object. You can also set some additional source properties. The PWX Bulk Metadata Adapter retrieves metadata for the source table from the PowerExchange-generated data map, including the field definitions.

In the user's environment, the target can reside on premises or in a cloud. If you do not want to send all of the incoming fields to the target, you can define one or more exclusion field rules. As part of configuring the target, you must also map the incoming fields to the target fields. You can map fields automatically based on the field names or map them manually.

VSAM source properties in mappings

When you create a mapping, you configure a Source transformation for a VSAM source and set some source properties for the Source transformation.

Click the Source transformation box on the mapping canvas. Then configure the source properties in the **Source Properties** panel below the mapping canvas.

If you have multiple Source transformations in a mapping, configure source properties for each source. If you use a specific connection instead of a connection parameter, enter the same connection for all sources.

Note: If a mapping includes source tables or columns that have special characters in their names, the associated mapping task fails because it cannot import the source metadata. Special characters include s #, $, @, %, *, !, and ~. Therefore, ensure that none of the source tables and mapped columns have names with any special character. Rename the tables and columns if necessary.

The **Source Properties** panel contains the following tabs on which you can enter information:

- **General** tab. Specify a user-defined name and optional description for the VSAM source.
- **Source** tab. Configure the source properties under **Details** and under **Advanced**. The advanced properties are available only if you enter a specific connection for the source. If you enter a parameter for the source connection, the advanced properties are not displayed. However, you can define these properties in the associated mapping task after you select a specific source connection for the task.

  Note: The Filter and Sort options under **Query Options** are not supported.

  The following table describes the VSAM source properties under **Details**:

Property	Description
Connection	*Required*. Select a source connection of the type **VSAM**, or click **New Connection** to create one. Alternatively, you can define a connection parameter in the mapping and enter a specific connection in each mapping task that is associated with the mapping. If you want to use a connection parameter with a specific source object, you must first select a specific connection so that you can access the source to select the source object in the **Object** field. After you select the source object, define the connection parameter. To define a connection parameter, click **New Parameter** next to the **Connection** property and enter the following information: - Enter a parameter name that begins with a letter and is up to 200 characters in length, including @, #, _, and alphanumeric characters. - In the **Type** property, verify that **connection** is selected. - Leave the **Connection Type** property blank. After you click **OK**, the parameter appears in the **Connection** property.
Source Type	*Required*. To use a parameter for the source object, select **Parameter**. To read data from a VSAM data source and write multiple record types to a VSAM target, select **Multi Group**.

Property	Description
Parameter	Parameter to use for the source object. This field only appears when you select **Parameter** as the source type. To parameterize a VSAM source table, click **New Parameter** to create a parameter. Enter a parameter name and then select the type **multi group data object**. This type of parameter can accept a set of groups. From the **New Input Parameter** dialog box, you can add or delete groups for the parameter. When you create a multi-group data object parameter, the corresponding links are also created. You can use these links in the mapping to connect to downstream transformations.
Schema, Data Map, and Tables	On the **Source** tab, click the pencil to select a schema, a data map, or table. For multiple record write, select **Use Sequence Fields (Select All Tables)** on the **Select Source Tables** dialog to preserve the data sequence as it passes to the target. When you select this option, all non-complex tables are automatically selected. Complex tables are not selected. **Note:** The **Preview Data** feature is not supported.

The following table describes the optional source properties under **Advanced**:

Advanced Property	Description
Filelist File	For VSAM files, select this check box only if you entered a filelist file in the **File Name Override** field. You must use a filelist file. A filelist file specifies the names of the files from which source data is to be retrieved. The PowerExchange Listener retrieves data from all of the files named in the filelist file in the order in which the files are listed. To request file list processing for a sequential or VSAM ESDS source, configure the following session attributes: - Specify the filelist file name in the **File Name Override** attribute. - Select the **Filelist File** option.
File Name Override	For VSAM files, overrides the data set or file name in the PowerExchange data map. Enter the complete data set or file name. If you select the **Filelist File** check box, enter the name of a filelist file in this attribute. A filelist file is a list of files.

Advanced Property	Description

Filter Overrides	One or more conditions that filter the source records that the PWX Bulk Reader retrieves from the PowerExchange Listener. Based on the filter conditions, a WHERE clause is appended to the default SQL SELECT query that the PWX Bulk Reader uses to read the records.
	VSAM Connector supports two forms of filter condition syntax. For single record source definitions such as single record nonrelational data maps, specify a single filter condition statement. You can specify a single filter condition or join numerous filter conditions by using the conditional operands that PowerExchange supports for NRDB SQL statements. For example:
	`column1 is NULL and column2='A'`
	For multi-record nonrelational source definitions, you can also use the following syntax:
	`group_name1=filter_condition;group_name2=filter_condition;...`
	Use the group_name form to specify filter conditions for one or more record types in a multi-record source definition. To apply a filter condition to all records in a multi-record source definition, use a single filter condition without group_name. You cannot combine single filter conditions with group_name conditions.
	You must use the nonrelational SQL syntax that PowerExchange supports for comparison and LIKE operators. For more information, see the "PowerExchange Nonrelational SQL" chapter in the *PowerExchange Reference Manual*. If the filter includes a date or timestamp column, ensure that the range of years in the column data is within the range of years for PowerExchange data filtering, as controlled by the DATERANGE statement in the DBMOVER configuration file. The default range is 1800 to 2200.
	For example, to select records in which the TYPE column has a value of A or D, specify the following condition:
	`TYPE='A' or TYPE='D'`
	Maximum length of the entire SELECT statement that PowerExchange builds, including the WHERE clause with the filter conditions, is 8192 bytes.
	Note: If you specify the **Filter Overrides** attribute and also specify an **SQL Query Override** attribute value that contains a filtering WHERE clause, the resulting SELECT statement contains a WHERE clause that uses the AND operator to associate the Filter Overrides filter conditions with the SQL Query Override conditions. For example:
	`SELECT * from schema.table WHERE Filter_Overrides_conditions AND SQL_Query_Override_conditions`

Flush After N Blocks	For multiple-record sources, specifies the maximum number of block flushes that can occur without any one block being flushed. For bulk multiple-record sources, by default, PWXPC flushes blocks of data only when the buffers are completely full or at end-of-file. If some record types do not have as much data as others, flushing might not occur often. In this case, the record types might not have data on the target for a long time, thereby blocking flushes on the writer side. To ensure that buffers for all record types are flushed at a regular interval, define this **Flush After N Blocks** session property. This property specifies the maximum number of block flushes that can occur across all record types without any one block being flushed. A value of zero disables this feature and causes flushing to occur only when blocks are full. Valid values for the property are -1 to 100000. The default value of -1 works in the following manner: For all multiple-record sources that do not use sequence fields, process the same as Flush After N Blocks = 0, which disables this feature and flushes only when blocks are full. For all multiple-record sources that use sequence fields, use Flush After N Blocks = 7 * (number of record types in the source).

Advanced Property	Description
Map Name Override	Overrides the data map name of the source PowerExchange data map.
Map Schema Override	Overrides the schema name in the source PowerExchange data map.
PWX Partition Strategy	For VSAM files, specifies one of the following partitioning strategies: **Single Connection**. PowerExchange creates a single connection to the data source. Any overrides specified for the first partition are used for all partitions. With this option, if you specify any overrides for other partitions that differ from the overrides for the first partition, the session fails with an error message. **Overrides Driven**. If the specified overrides are the same for all partitions, PowerExchange creates a single connection to the data source. If the overrides are not identical for all partitions, PowerExchange creates multiple connections.

SQL Query Override	An SQL statement that overrides the default SQL query that the PWX Bulk Reader uses to retrieve records from PowerExchange.
	The connector replaces the default SQL query with the SQL statement that you enter and passes the SQL statement to Cloud Data Integration for processing.
	You must use the nonrelational SQL syntax that PowerExchange supports. For more information, see the "PowerExchange Nonrelational SQL" chapter in the *PowerExchange Reference Manual*.
	If the override statement includes a date or timestamp column, ensure that the range of years in the column data is within the range of years for PowerExchange data checking and filtering, as controlled by the DATERANGE statement in the DBMOVER configuration file. The default range is 1800 to 2200.
	For example, to select records for the USER source table when the TYPE column has the value of A or D, specify the following query:
	`SELECT ID, NAME from USER where TYPE='A' or TYPE='D';`
	Maximum length of the SELECT query override statement is 8192 bytes.
	Note: If you specify an **SQL Query Override** attribute value that contains a filtering WHERE clause and also specify the **Filter Overrides** attribute, the resulting SELECT statement contains a WHERE clause that uses the AND operator to associate the SQL Query Override conditions with the Filter Overrides filter conditions. For example:
	`SELECT * from schema.table WHERE Filter_Overrides_conditions AND SQL_Query_Override_conditions`
	For a multiple-record source, use the following syntax:
	`group_name1=sql_query_override1; group_name2=sql_query_override2;...`
	For example, you can select only records with ID column values that contain DBA for a multi-record source with two records called USER1 and USER2 by specifying the following SQL query override:
	`USER1=Select ID, NAME from USER1 where ID='DBA'; USER2=Select ID, NAME from USER2 where ID='DBA';`

Tracing Level	The level of detail that appears in the session log for a Source transformation. Select one of the following options, which are listed in order of increasing detail: - **Terse** - **Normal** - **Verbose Initialization** - **Verbose Data** Default is **Normal**.

- **Fields** tab. Select the fields that you want to use in the data flow. You can edit field metadata such as precision and scale if necessary.

Note: Ignore the **Partitions** tab. Key-range partitioning is not supported.

Targets in mappings with VSAM sources

To configure a target in a mapping, click the Target transformation box on the mapping canvas. Then configure the target properties in the **Target Properties** panel below the mapping canvas.

Note: To determine which target types are supported by the VSAM Connector, see "Supported targets" on page 11.

- **General** tab. Specify a user-defined name and optional description for the target.
- **Incoming Fields** tab. Optionally review the included fields and define field exclusion rules and renaming criteria.
- **Target** tab. Configure the target properties under **Details** and under **Advanced**. The advanced properties are available only if you enter a specific connection for the target. Target properties vary based on the supported target that you select.
- **Target Fields** tab. Optionally add fields, delete fields, or edit field metadata such as native data type, precision, and scale.
- **Field Mapping** tab. Map the incoming fields to the target fields. You can manually map an incoming field to a target field or automatically map fields based on the field names. If you entered a connection parameter, you must select **Automatic** as the field map option to be able to validate the mapping. Alternatively, you can completely or partially parameterize field mappings, and then define the specific field mappings in the mapping task. For more information, see Transformations > Target Transformation > Target Transformation Field Mappings in the Data Integration Help.

For more information, see the appropriate target connector guide for target for specific instructions.

VSAM targets in a mapping

To write data to a VSAM data set, configure a VSAM object as the Target transformation in a mapping.

Specify the name and description of the VSAM target. Configure the target and advanced properties for the target object.

The following table describes the target properties under **Details**:

Property	Description
Connection	Select an existing VSAM target connection, or create a target connection that has a type of **Adabas**. Alternatively, you can use a connection parameter in the mapping and then enter a specific connection in each mapping task that is associated with the mapping. If you want to use a connection parameter with a specific target object, you must first select a specific connection so that you can access the target to select the target object in the **Object** field. After you select the target object, define the connection parameter. If you want to use parameters for both the connection and target object, you do not need to first select a specific connection. You can just define the parameters in either order. To define a connection parameter, click **New Parameter** and enter the following information: - Enter a parameter name that begins with a letter and is up to 200 characters in length, including @, #, _, and alphanumeric characters. - In the **Type** property, verify that **connection** is selected. - Leave the **Connection Type** property blank.
Target Type	To specify a target table, select **Single Object**. To write multiple record types to a VSAM target, select **Multi Group**. If you selected **Multi Group** and **Use Sequence Fields (Select All Tables)** for the source, make sure **Use Sequence Fields (Select All Tables)** is also selected for the target. This setting is required for multiple-record write. **Note:** - For multiple-record writes, you must join source and target tables and manually map the source fields to the target fields through the **Target Field Group** field. - For multiple-record writes, if you select the same data map as a target, you must override the file name to avoid overwriting the source file.
Object	Click **Select**. In the **Target Object** dialog box, select **Existing** and then select an existing target table.
Operation	Select the target operation. You can perform the following operations on an VSAM target: - Insert - Update - Upsert - Delete - Data Driven

The following table describes the advanced target properties that appear under **Advanced**:

Advanced Property	Description
Forward Rejected Rows	Select this option to have the Target transformation pass rejected records to a reject file that is generated in the following directory: `\agent_install_directory\apps\Data_Integration_Server\data\error` Clear this option to ignore the rejected records. **Note:** This field is not displayed if a parameter is specified for the target connection.
Pre SQL	SQL statements that run on the target before the extracted data is written to a target. Maximum length is 5000 characters.
Post SQL	SQL statements that run on the target after the extracted data is written to a target. Maximum length is 5000 characters.

Advanced Property	Description
Update SQL Override	Overrides the default Update SQL that is sent to PowerExchange.
Map Schema Override	A schema name that overrides the schema name in the PowerExchange data map.
Map Name Override	Overrides the target PowerExchange data map name.
Insert SQL override	Overrides the default Insert SQL sent to PowerExchange.
Delete SQL override	Overrides the default Delete SQL that is sent to PowerExchange.
Pre SQL run once per Connection	Runs the SQL that you specify in the Pre SQL attribute only once for a connection. Select this attribute in either of the following cases: - In the **Pre SQL** attribute for a session that uses writer partitioning, you specify a SQL statement such as CREATEFILE that can run only once for the session. If you do not select **Pre SQL run once per Connection**, the session tries to run the statement once for each partition. - In the **Pre SQL** attribute for a session that performs a multiple-record write, you specify a CREATEFILE statement that creates a new generation of a GDG or creates an empty file. If you do not select **Pre SQL run once per Connection**, the session creates a generation or tries to create a new empty file for each record that the session writes.
Truncate Target	Truncates, or deletes, table contents before loading new data. **Note:** VSAM data sets must be defined with the REUSE option for this truncate option to function correctly.
Initialize Target	Select this option to have PowerExchange allow both inserts and updates into empty VSAM data sets. If this option is not selected, PowerExchange only allows inserts into empty VSAM data sets.
File Name Override	Overrides the data set in the PowerExchange data map. Enter the complete data set name.
PWX Partition Strategy	Specify one of the following partitioning strategies: - **Single Connection**. PowerExchange creates a single connection to the data source. Any overrides specified for the first partition are used for all partitions. With this option, if you specify any overrides for other partitions that differ from the overrides for the first partition, the session fails with an error message. - **Overrides Driven**. If the specified overrides are the same for all partitions, PowerExchange creates a single connection to the data source. If the overrides are not identical for all partitions, PowerExchange creates multiple connections.

.

Lookup Transformation

You can configure a cached Lookup transformation when you use a VSAM connection in a mapping to return data from a VSAM source based on a specified lookup condition. However, the VSAM Connector does not support uncached lookup.

For more information, see Cloud Data Integration *Transformations*.

Multiple-record writes to nonrelational targets

You can use a multiple-record data map to read data from a nonrelational data source and write the multiple record types in a single pass to a nonrelational target. This process is called a multiple-record write.

When Cloud Data Integration performs a multiple-record write, it preserves source sequencing information. To enable multiple-record writes with sequencing, select the **Use Sequence Fields (Select All Tables)** option in the **Select Source Tables** dialog box for both the source and target definition. When you select these options, Cloud Data Integration generates primary key (DTL_PK) and foreign key fields (DTL_FK) for each group in the source and target definitions, which enables Cloud Data Integration to reassemble the data in the correct sequence at the target.

You can perform multiple-record writes with the following types of sources and targets:

- Source and target IMS unload and VSAM data sets on z/OS
- Any combination of the following types of nonrelational sources and targets:
  - Files on Linux, UNIX, or Windows
  - Files on IBM i

Mapping task configuration

Configure a mapping task for a mapping with a VSAM source. When you run the mapping task, the VSAM records are transmitted to the target.

Note: SQL ELT optimization and Visio templates are not supported for VSAM mapping tasks.

To begin configuring a mapping task, perform one of the following actions:

- In Data Integration, click **New** > **Mapping Task** and then click **Create**.
- While the mapping is open, select **New Mapping Task** from the **Actions** menu.

Then enter information on some or all of the following pages in the mapping task wizard:

- **General**. Enter a name for the mapping task, a project or folder location for the mapping task, and the runtime environment that contains the Secure Agent or Secure Agent Group under which the mapping task will run. If you began creating the mapping task without first opening the mapping, you must also select the mapping with which to associate the mapping task. The description is optional.
- **Sources**. This page appears only if you defined a parameter for the connection. If you defined a connection parameter, select or create the specific connection that you want the mapping task to use.
- **Targets**. This page appears only if you defined a parameter for the target connection or object in the mapping. For more information about configuring targets in mapping tasks, click the help icon on this page.
- **Runtime Options**. Optionally define a schedule for running the mapping task.

Properties for parameterized sources and connections

The **Sources** page appears in the mapping task wizard if you defined a parameter for the connection or source object in the associated mapping.

If you defined a parameter for the source connection in the mapping, you must select or create a specific connection for the mapping task. Ensure that the connection type is consistent with the source object type. After you select a connection, you can click **Advanced Options** to configure the optional advanced properties.

For more information about these advanced properties, see "VSAM source properties in mappings" on page 18.

Note: If the mapping uses a specific source connection, you can specify these advanced properties in the mapping and also specify them in the mapping task. In this case, the properties in the mapping task override those in the mapping.

If you are parameterizing a source table, select the tables that correspond to the groups you selected in the mapping from the mapping task **Sources** tab. The number of tables you select must match the number of groups that were added in the multi group data object parameter. To select the source tables, click **Select**.

Note: To parameterize a source table, you must define a parameter of the type **multi group data object** for the mapping.

By default, the tables are associated with the groups in the order in which the tables appear in the table selection window. From the **Group Map** area, you can change the group to table association by using drag and drop to associate a group with a table.

When you finish entering information on the **Sources** page, click **Validate** to validate your entries.

Note: If the mapping specifies a connection parameter and you edit the mapping to change the source object, check the specific connection in the mapping task to determine if you need to reset it and then **Validate** the mapping task again.

Running a mapping task

A mapping task runs on-premises under control of the Secure Agent.

To run the mapping task, use any of the following methods:

- After you finish configuring the mapping task, click **Run**.
- Manually start the mapping task from the **Explore** page.
- Create a schedule that runs the mapping task based on a date, time, and repeat frequency, such as hourly, daily, weekly, or monthly.
- Include the mapping task in a taskflow and then run the taskflow.

Monitoring a mapping task

After you run a mapping task, you can check the job status for the task to determine if the job succeeded, completed with warnings, or failed. If the job failed or ended with a warning, you can view the session log for the job to identify the error or warning conditions.

Each time a user or schedule starts a mapping task, Informatica Intelligent Cloud Services creates a job to run the task.

1. To view a list of mapping task jobs, use one of the following methods:
    - In Data Integration, select **My Jobs**. The **My Jobs** page lists all jobs that you, the user who is currently logged in, started. It does not include jobs that were started by a schedule.

- In Monitor, select **All Jobs** or **Running Jobs**. The **All Jobs** page lists all of the running and completed jobs that members of your organization started or that were started by a schedule. The **Running Jobs** page lists only the jobs that are currently running or that have completed within the last 5 minutes.

2. On the **My Jobs**, **All Jobs**, or **Running Jobs** page, review the information about the jobs you're interested in. The following information is displayed by default and shown for each job instance:
   - Start date and time
   - End date and time
   - Number of rows processed
   - Current state, which can be one of the following values: Queued, Starting, Running, Success, Stopped, Suspended, Warning, or Failed.

   Tip: You can change the columns that are displayed by right-clicking the column-heading area and selecting or deselecting column names in the list.

   For more information about a job, click the job name to open a page that contains job details and a download link for the session log. You can also download the session log by moving the cursor to the right end of the row for a job and clicking the **Download Log** icon.
3. If the job state is Failed or Warning, gather more information about the error or warning condition in the following ways:
   - For a Failed job, hover the cursor over the Failed icon in the **State** column. A brief description of the error condition appears, which you can copy.
   - Download the session log and look for error or warning messages that begin with the message ID prefix of "PWXPC" or "PWX."

   If possible, correct the error or warning condition and then restart the job. To restart the job from the **Jobs** or **My Jobs** page, move the cursor to the right end of the row for the job and click the **Restart** icon.

CHAPTER 4

VSAM administration

As the organization administrator of a VSAM Connector implementation, you might need to perform some administrative tasks to keep processing of records running smoothly.

To learn about some common administrative issues and tasks, read the following topic:

- "Upgrading to a new version of PowerExchange" on page 28

Upgrading to a new version of PowerExchange

The connector can retrieve data from only a single version of PowerExchange. If you upgrade your on-premise environment to a later supported PowerExchange version, use this task to move your Informatica Intelligent Cloud Services Data Integration connections to the new PowerExchange environment.

1. Upgrade the PowerExchange installation. For more information, see the *PowerExchange Installation and Upgrade Guide*.
2. In Monitor, open the **Running Jobs** page.
3. Make sure that any running mapping tasks that connect to the previous PowerExchange version run to completion, or stop the task if necessary.
4. Stop the Secure Agent that is specified in the connection configuration for any connection to the previous PowerExchange version.
5. In Data Integration, edit the connections to point to the new PowerExchange version.
6. If necessary, edit the mapping task configurations to specify the appropriate connection to the new version of the PowerExchange environment.
7. Start the mapping tasks again, when appropriate.

CHAPTER 5

Data type reference

Data Integration uses the following data types in mappings and tasks with VSAM:

VSAM native data types

VSAM data types appear in the source and target transformations when you choose to edit metadata for the fields.

Transformation data types

Set of data types that appear in the transformations. They are internal data types based on ANSI SQL-92 generic data types, which the Secure Agent uses to move data across platforms. Transformation data types appear in all transformations in a mapping.

VSAM and transformation data types

When Data Integration reads source data, it converts the native data types to the comparable transformation data types before transforming the data. When Data Integration writes to a target, it converts the transformation data types to the comparable native data types.

The following table lists the VSAM data types that Data Integration supports and the corresponding transformation data types:

PowerExchange Datatype	Precision	Transformation Datatype	Range
BIN	10	Binary	1 to 104,857,600 bytes You can pass binary data from a source to a target, but you cannot perform transformations on binary data. PowerCenter does not support binary data for COBOL or flat file sources.
CHAR	10	String	1 to 104,857,600 characters Fixed-length or varying-length string.
DATE	10	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. Combined date/time value, with precision to the nanosecond.
DOUBLE	18	Double	Precision of 15 digits Double-precision floating-point numeric value.
FLOAT	7	Double	Precision of 15 digits Double-precision floating-point numeric value.

PowerExchange Datatype	Precision	Transformation Datatype	Range
NUM8	3	Small Integer	Precision of 5 and scale of 0 Integer value.
NUM8U	3	Small Integer	Precision of 5 and scale of 0 Integer value.
NUM16	5	Small Integer	Precision of 5 and scale of 0 Integer value.
NUM16U	5	Integer	Precision of 10 and scale of 0 Integer value.
NUM32	10	Integer	Precision of 10 and scale of 0 Integer value.
NUM32U	10	Double	Precision of 15 digits Double-precision floating-point numeric value.
NUM64	19	Decimal	Precision 1 to 28 digits, scale 0 to 28 Decimal value with declared precision and scale. Scale must be less than or equal to precision. If you pass a value with negative scale or declared precision greater than 28, the PowerCenter Integration Service converts it to a double.
NUM64U	19	Decimal	Precision 1 to 28 digits, scale 0 to 28 Decimal value with declared precision and scale. Scale must be less than or equal to precision. If you pass a value with negative scale or declared precision greater than 28, the PowerCenter Integration Service converts it to a double.
NUMCHAR	-	String	1 to 104,857,600 characters Fixed-length or varying-length string.
PACKED	15	Decimal	Precision 1 to 28 digits, scale 0 to 28 Decimal value with declared precision and scale. Scale must be less than or equal to precision. If you pass a value with negative scale or declared precision greater than 28, the PowerCenter Integration Service converts it to a double.
TIME	5	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. Combined date/time value, with precision to the nanosecond.
TIMESTAMP	5	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. Combined date/time value, with precision to the nanosecond.
UPACKED	15	Decimal	Precision 1 to 28 digits, scale 0 to 28 Decimal value with declared precision and scale. Scale must be less than or equal to precision. If you pass a value with negative scale or declared precision greater than 28, the PowerCenter Integration Service converts it to a double.

PowerExchange Datatype	Precision	Transformation Datatype	Range
UZONED	15	Decimal	Precision 1 to 28 digits, scale 0 to 28 Decimal value with declared precision and scale. Scale must be less than or equal to precision. If you pass a value with negative scale or declared precision greater than 28, the PowerCenter Integration Service converts it to a double.
VARBIN	10	Binary	1 to 104,857,600 bytes You can pass binary data from a source to a target, but you cannot perform transformations on binary data. PowerCenter does not support binary data for COBOL or flat file sources.
VARCHAR	10	String	1 to 104,857,600 characters Fixed-length or varying-length string.
ZONED	15	Decimal	Precision 1 to 28 digits, scale 0 to 28 Decimal value with declared precision and scale. Scale must be less than or equal to precision. If you pass a value with negative scale or declared precision greater than 28, the PowerCenter Integration Service converts it to a double.

CHAPTER 6

Connection overrides reference

When you create a mapping that uses this connector type, you can specify a connection property that allows you to specify connection property overrides through a parameter file in the connection properties panel **Custom Properties** field.

Connection overrides

The following table lists the connection overrides that you can specify for nonrelational sources such as Adabas, Datacom, IDMS, IMS, sequential file, and VSAM connections:

Connection Override	Default	Valid Values
Array Size	25	1 to 100000
Connection Retry Period	0	0 to 1000000
Offload Processing	No	A - Auto F - Filter After N - No Y - Yes
Offload Threads	0	0 to 64
Password	no default	A valid password for the specified user
Replace Low-values With Spaces	no default	Y|N
User Name	no default	A user name
Confirm Network Write	Confirm Write On	Y|N

INDEX

C

Cloud Application Integration community
  URL 5
Cloud Developer community
  URL 5
connection
  overrides 32
connection overrides reference
  overview 32
connections
  VSAM 14
  VSAM connection properties 14

D

Data Integration community
  URL 5
data type reference
  overview 29

I

Informatica Global Customer Support
  contact information 6
Informatica Intelligent Cloud Services
  web site 5

L

Lookup transformation 24

M

maintenance outages 6
mapping tasks
  monitoring mapping tasks 26
  running mapping tasks 26
  VSAM 25
mappings
  VSAM 17
Multiple-Record writes 25

S

status
  Informatica Intelligent Cloud Services 6
system status 6

T

trust site
  description 6

U

upgrade notifications 6

V

VSAM
  mapping task configuration 25
  source properties in mappings 18
  administration 28
  connection properties 14
  connections overview 14
  Connector architecture 9
  Connector software requirements 11
  Connector usage scenario 7
  data types 29
  introduction to Connector 7
  mapping configuration 17
  mapping task properties for parameterized sources 25
  prerequisite tasks for using the Connector 12
  rules and guidelines for Connector use 12
  supported target types 11
  target properties in VSAM source mappings 22
  task flow for using the Connector 13
VSAM Connector
  task and asset types 10
VSAM targets
  mapping 22

W

web site 5