Informatica

Informatica® Cloud Data Integration

Microsoft Azure SQL Data Warehouse V2 Connector

Informatica Cloud Data Integration Microsoft Azure SQL Data Warehouse V2 Connector
October 2020

© Copyright Informatica LLC 2017, 2023

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: 2023-11-29

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 Microsoft Azure SQL Data Warehouse V2 Connector 7

Microsoft Azure SQL Data Warehouse V2 Connector Overview 7

Microsoft Azure SQL Data Warehouse V2 Connector Task and Object Types 8

Administration of Microsoft Azure SQL Data Warehouse V2 Connector 8

- Azure Active Directory Authentication 8

Chapter 2: Microsoft Azure SQL Data Warehouse V2 Connections 10

Microsoft Azure SQL Data Warehouse V2 Connection Overview 10

Microsoft Azure SQL Data Warehouse V2 connection properties 11

Chapter 3: Synchronization Tasks with Microsoft Azure SQL Data Warehouse V2 12

IDENTITY Column 12

Microsoft Azure SQL Data Warehouse V2 Source Properties 12

Microsoft Azure SQL Data Warehouse V2 Lookups in Synchronization Tasks 14

Microsoft Azure SQL Data Warehouse V2 Target Properties 14

Rules and guidelines for synchronization tasks 15

Chapter 4: Mapping Tasks with Microsoft Azure SQL Data Warehouse V2 16

Microsoft Azure SQL Data Warehouse V2 Source Properties 16

- Key Range Partitioning 17
- Configuring Key Range Partitioning 18

Microsoft Azure SQL Data Warehouse V2 Lookups in Mapping Configuration Tasks 18

Microsoft Azure SQL Data Warehouse V2 Target Properties in Mappings and Mapping Configuration Tasks 19

Chapter 5: SQL ELT Optimization 21

SQL ELT Optimization Overview 21

Install the ODBC Driver 21

Create a System DSN. 22
Create an ODBC Connection. 25
Cross-Schema SQL ELT Optimization. 25
SQL ELT Optimization Functions. 26
Rules and Guidelines for Functions in SQL ELT Optimization. 29

Chapter 6: Data Type Reference. 32
Microsoft Azure SQL Data Warehouse V2 Data Type Reference Overview. 32
Microsoft Azure SQL Data Warehouse V2 and Transformation Data Types. 33

Index. 34

Preface

Use *Microsoft Azure SQL Data Warehouse V2 Connector* to learn how to read from or write to Microsoft Azure SQL Data Warehouse by using Cloud Data Integration. Learn to create a connection, develop and run synchronization tasks, mappings, and mapping tasks. You can also learn how to configure pushdown optimization using an ODBC connection.

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

To get the latest documentation for your product, browse the Informatica Knowledge Base at https://kb.informatica.com/_layouts/ProductDocumentation/Page/ProductDocumentSearch.aspx.

If you have questions, comments, or ideas about this documentation, contact the Informatica Documentation team through email 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 Data Integration editions and applications as well as information about other 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/community/collections/cloud_integration

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, go to the Informatica Intelligent Cloud Services Status page and click **SUBSCRIBE TO UPDATES**. You can then 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 Customer Support Center by telephone or online.

For online support, click **Submit Support Request** in Informatica Intelligent Cloud Services. You can also use Online Support to log a case. Online Support requires a login. You can request a login at https://network.informatica.com/welcome.

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 Microsoft Azure SQL Data Warehouse V2 Connector

This chapter includes the following topics:

- Microsoft Azure SQL Data Warehouse V2 Connector Overview, 7
- Microsoft Azure SQL Data Warehouse V2 Connector Task and Object Types, 8
- Administration of Microsoft Azure SQL Data Warehouse V2 Connector, 8

Microsoft Azure SQL Data Warehouse V2 Connector Overview

Microsoft Azure SQL Data Warehouse V2 Connector enables you to read data from or write data to Microsoft Azure SQL Data Warehouse. You can insert, upsert, or update data to or delete data from a Microsoft Azure SQL Data Warehouse V2 target.

Microsoft Azure SQL Data Warehouse is an enterprise-class, data-warehouse-as-a-service. It is a distributed database in the Azure Cloud that can process relational and non-relational data.

When you read data from or write data to Microsoft Azure SQL Data Warehouse, you can specify the Hosted Agent or the Secure Agent.

Microsoft Azure SQL Data Warehouse V2 Connector is optimized for large data sets and can perform better than traditional data integration methods, such as ODBC or JDBC. When you read data from or write data to a Microsoft Azure SQL Data Warehouse, the connector stages data files to Microsoft Azure Storage and uses T-SQL commands with Microsoft Polybase to extract or load relational and non-relational data in parallel. You can use Microsoft SQL Server or Azure Active Directory authentication to connect to a Microsoft Azure SQL Data Warehouse.

For example, you work in sales operations and you frequently need to analyze a high volume of data to improve operational intelligence. You design a mapping to read data from Salesforce and other transactional systems and aggregate the data. You create a summary table in Microsoft Azure SQL Data Warehouse that you can query against to assess your sales organization's performance.

Note: When upgrading from Microsoft Azure SQL Data Warehouse Connector to Microsoft Azure SQL Data Warehouse V2 Connector, you must reimport the source and target definitions using the Azure SQL Data Warehouse V2 connection in mappings and tasks.

Microsoft Azure SQL Data Warehouse V2 Connector Task and Object Types

The following table lists the Microsoft Azure SQL Data Warehouse V2 object types that you can use in Data Integration tasks:

Task Type	Source	Target	Lookup
Synchronization	Yes	Yes	Yes
Mapping	Yes	Yes	Yes

Administration of Microsoft Azure SQL Data Warehouse V2 Connector

As a user, you can use Microsoft Azure SQL Data Warehouse V2 Connector after the organization administrator performs the following tasks:

- Obtains the JDBC URL from Microsoft Azure SQL Data Warehouse.
- Verify that either the `db_owner` privilege or the following more granular privileges are granted to the user to connect to the Microsoft Azure SQL Data Warehouse and perform read and write operations successfully:
  - `EXEC sp_addrolemember 'db_datareader', '<user>';` // Alternately assign permission to individual table
  - `EXEC sp_addrolemember 'db_datawriter', '<user>';` // Alternately assign permission to individual table
  - `GRANT ALTER ANY EXTERNAL DATA SOURCE TO <user>;`
  - `GRANT ALTER ANY EXTERNAL FILE FORMAT TO <user>;`
  - `GRANT CONTROL TO <user>;`
  - `GRANT CREATE TABLE TO <user>;`
  - Assign required privileges for tasks performed through Pre-SQL and Post-SQL commands.

Azure Active Directory Authentication

You can configure Azure Active Directory (AAD) authentication to connect to Microsoft Azure SQL Data Warehouse. Default is Microsoft SQL Server authentication.

To configure the AAD authentication, perform the following tasks:

Import Server Certificate

Applicable if you want to use AAD authentication with Active Directory Federation Services in Azure. Import the server certificate to the `<Secure Agent installation directory>\jre\lib\security\cacerts` file. Use the following command:

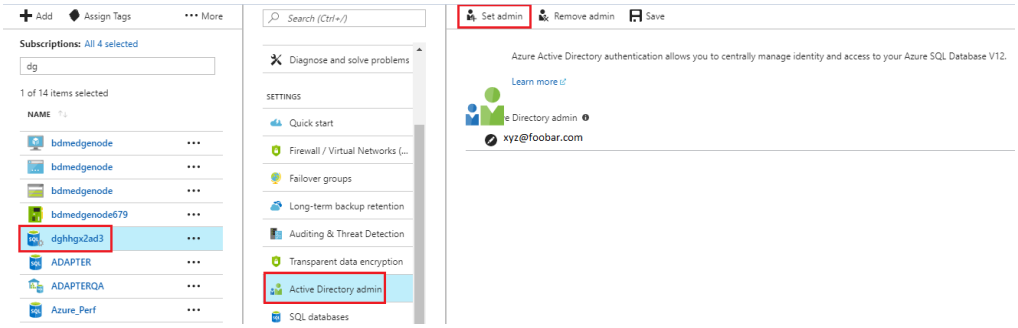
```
keytool -import -trustcacerts -alias <alias name of the certificate> -file <certificate file path> -keystore <Secure Agent installation directory>\jre\lib\security\cacerts -storepass <password for the truststore>
```

Note: You cannot run a task on a Hosted Agent, if Active Directory Federation Services is used for AAD authentication.

Set Admin

Perform the following steps to set admin between Microsoft SQL Server that has the Microsoft Azure SQL Data Warehouse hosted and the Azure Active Directory:

1. Log on to the Microsoft Azure portal using your credentials. The Dashboard page appears.
2. From the All Resources page, select the Microsoft SQL Server that has the Microsoft Azure SQL Data Warehouse hosted.
3. Select the **Active Directory admin** option under Settings displayed for the Microsoft SQL Server. The image shows the Active Directory admin settings:



4. Click **Set admin**. The Add admin page appears.
5. Enter the email ID that you want to use as admin and click **Select**.
6. Click **Save**.

Create a User

Perform the following steps to create a user:

1. Connect to the Microsoft Azure SQL Data Warehouse using the Azure Active Directory admin created in the previous steps. You can use Microsoft SQL Server Management Studio to connect to the Microsoft Azure SQL Data Warehouse.
2. Type and run the following command to create a user: `create user [user@foobar.com] from external provider;`
   For more information, see Microsoft Azure documentation.
3. Assign the required privileges to the user.

Configure the JDBC URL and the user you created in connection properties to enable AAD authentication.

CHAPTER 2

Microsoft Azure SQL Data Warehouse V2 Connections

This chapter includes the following topics:

- Microsoft Azure SQL Data Warehouse V2 Connection Overview, 10
- Microsoft Azure SQL Data Warehouse V2 connection properties, 11

Microsoft Azure SQL Data Warehouse V2 Connection Overview

Create a Microsoft Azure SQL Data Warehouse V2 connection to access Microsoft Azure SQL Data Warehouse data from Data Integration. You can create a connection on the **Connections** page or when you create a mapping or a task. After you create a connection, it becomes available to all users who have access to the organization.

Microsoft Azure SQL Data Warehouse V2 connection properties

The following table describes Microsoft Azure SQL Data Warehouse V2 connection properties:

Connection property	Description
Runtime Environment	The name of the runtime environment where you want to run the tasks.

Azure DW JDBC URL	Microsoft Azure Data Warehouse JDBC connection string.
	Example for Microsoft SQL Server authentication:
	`jdbc:sqlserver://<Server>.database.windows.net:1433;database=<Database>`
	Example for Azure Active Directory (AAD) authentication:
	`jdbc:sqlserver://<Server>.database.windows.net:1433;database=<Database>;encrypt=true;trustServerCertificate=false;hostNameInCertificate=*.database.windows.net;loginTimeout=30;Authentication=ActiveDirectoryPassword;`
	The default authentication is Microsoft SQL Server authentication.

Azure DW JDBC Username	User name to connect to the Microsoft Azure SQL Data Warehouse account. Provide AAD user name for AAD authentication.
Azure DW JDBC Password	Password to connect to the Microsoft Azure SQL Data Warehouse account.
Azure DW Schema Name	Name of the schema in Microsoft Azure SQL Data Warehouse.
Azure Blob Account Name	Name of the Microsoft Azure Storage account to stage the files.
Azure Blob Account Key	Microsoft Azure Storage access key to stage the files.

CHAPTER 3

Synchronization Tasks with Microsoft Azure SQL Data Warehouse V2

The Synchronization application allows you to synchronize data between a source and target.

You can configure a synchronization task using the Synchronization Task wizard. You can use Microsoft Azure SQL Data Warehouse objects as sources, targets, or lookup objects. You can use expressions to transform the data according to your business logic, use data filters to filter data before writing it to targets, sort data in ascending or descending order of multiple fields.

When you create a task, you can associate it with a schedule to run it at specified times or on regular intervals. Or, you can run it manually. You can monitor tasks that are currently running in the activity monitor and view logs about completed tasks in the activity log.

IDENTITY Column

You can use the IDENTITY column in synchronization and mapping tasks. Consider the following instructions when you use the IDENTITY column:

- You can map an IDENTITY column for the source and lookup objects.
- You cannot map the IDENTITY column for target objects.
- You cannot define an IDENTITY column when you use Create Target.
- The Connector does not highlight the IDENTITY column in the imported objects.

Microsoft Azure SQL Data Warehouse V2 Source Properties

You can use a Microsoft Azure SQL Data Warehouse object as a single source in a synchronization task. You can configure Microsoft Azure SQL Data Warehouse V2 source properties on the **Source** page of the Synchronization Task wizard.

The following table describes the source properties for Microsoft Azure SQL Data Warehouse V2 Connector:

Source Property	Description
Connection	Name of the source connection.
Source Type	Type of the source object. Select Single.
Source Object	Name of the source object.
Display technical names instead of labels	Displays technical names instead of business names.
Display source fields in alphabetical order	Displays source fields in alphabetical order instead of the order returned by the source system.

Configure the Microsoft Azure SQL Data Warehouse V2 advanced source properties on the **Schedule** page of the Synchronization Task wizard.

The following table describes the advanced source properties for Microsoft Azure SQL Data Warehouse V2 Connector:

Advanced Source Property	Description
Azure Blob Container Name	Name of the container in Microsoft Azure Storage to use for staging before extracting data from Microsoft Azure SQL Data Warehouse.
Field Delimiter	Character used to separate fields in the file. Default is 0x1e. You can select the following field delimiters from the list: `~ ` ` | . TAB 0x1e`
Number of Concurrent Connections to Blob Store	Number of concurrent connections to extract data to the staging area in Microsoft Azure Blob Storage. Default is 4.
Quote Character	The Secure Agent Skips the specified character when you read data from Microsoft Azure SQL Data Warehouse.
SQL Override	When you read data from a Microsoft Azure SQL Data Warehouse object, you can configure SQL overrides and define constraints.
Pre-SQL	Pre-SQL command that must be run before reading data from the source.
Post- SQL	Post-SQL command that must be run after writing data to the target.

Microsoft Azure SQL Data Warehouse V2 Lookups in Synchronization Tasks

When you configure field mappings in a synchronization task, you can create a lookup to a Microsoft Azure SQL Data Warehouse object.

When you use a Microsoft Azure SQL Data Warehouse object as a lookup, you need to specify the Azure Blob Container Name, Field Delimiter, and Number of Concurrent Connections to Blob Store properties in the Lookup tab. When you run a lookup task, the Secure Agent does not consider the specified values of the Azure Blob Container Name, Field Delimiter, and Number of Concurrent Connections to Blob Store properties in the Lookup tab.

You cannot use Binary, Varbinary, and Datetimeoffset datatypes columns as lookup conditions.

Microsoft Azure SQL Data Warehouse V2 Target Properties

You can use a Microsoft Azure SQL Data Warehouse object as a target in a synchronization task. When you use Microsoft Azure SQL Data Warehouse target objects, you can select a standard object as the primary target. You can configure Microsoft Azure SQL Data Warehouse target properties on the **Target** page of the Synchronization Task wizard.

The following table describes the target properties for Microsoft Azure SQL Data Warehouse V2:

Target Property	Description
Connection	Name of the target connection.
Target Object	Name of the target object. Select the primary target object.
Create Target	Creates a target. Enter a name for the target object and select the source fields that you want to use. By default, all source fields are used.
Display technical names instead of labels	Displays technical names instead of business names.
Display source fields in alphabetical order	Displays source fields in alphabetical order instead of the order returned by the source system.

Configure the Microsoft Azure SQL Data Warehouse V2 advanced target properties on the **Schedule** page of the Synchronization Task wizard.

The following table describes the advanced target properties for Microsoft Azure SQL Data Warehouse V2:

Advanced Target Property	Description
Azure Blob Container Name	Name of the container in Microsoft Azure Storage to use for staging before writing to Microsoft Azure SQL Data Warehouse.
Field Delimiter	Character used to separate fields in the file. Default is 0x1e. You can select the following field delimiters from the list: `~ ` ` | . TAB 0x1e`
Update Column	The key columns to upsert or update data to or delete data from Microsoft Azure SQL Data Warehouse.
Number of Concurrent Connections to Blob Store	Number of threads to use to move data to the staging area in Microsoft Azure Blob Storage. Default is 4.
Truncate Table	Truncates the target data before inserting the data to the target. This option is applicable for insert operation.
Quote Character	Specifies the quote character to skip when you write data to Microsoft Azure SQL Data Warehouse. When you write data to Microsoft Azure SQL Data Warehouse and the source table contains the specified quote character, the task fails. Change the quote character value to a value that does not exist in the source table.
Pre-SQL	Pre-SQL command that must be run before reading data from the source.
Post- SQL	Post-SQL command that must be run after writing data to the target.
Success File Directory	NA
Error File Directory	NA

Rules and guidelines for synchronization tasks

Consider the following rules and guidelines when you configure a synchronization task:

- When you use a lookup condition on a Microsoft Azure SQL Data Warehouse object and enable report error on multiple matched fields, the value is returned for any random field.
- You cannot schedule a task to run later if you use a lookup condition on a Microsoft Azure SQL Data Warehouse object.

CHAPTER 4

Mapping Tasks with Microsoft Azure SQL Data Warehouse V2

Use a mapping to define data flow logic that is not available in synchronization tasks, such as specific ordering of logic or joining sources from different systems. Use the Data Integration Mapping Designer to configure mappings.

When you configure a mapping to describe the flow of data from source and target, you can also add transformations to transform data. A transformation includes field rules to define incoming fields. Links visually represent how data moves through the data flow.

After you create a mapping, you can run the mapping or you can deploy the mapping in a mapping task. The mapping task allows you to process data based on the data flow logic defined in a mapping.

Use the Mapping task wizard to create a mapping task. When you create a mapping task, you select the mapping for the task to use.

If you configured parameters, which are placeholders for information, in a mapping, you can define the parameters in the mapping task. Defining parameters provides additional flexibility and allows you to use the same mapping in multiple mapping tasks. For example, you can use a parameter for a source connection in a mapping, and then define the source connection when you configure the mapping task.

When you create a mapping task, you can associate the task with a schedule to run it at specified times or on regular intervals. Or, you can run it manually. You can also configure advanced session properties. You can monitor tasks that are currently running in the activity monitor and view details about completed tasks in the activity log.

Microsoft Azure SQL Data Warehouse V2 Source Properties

In a mapping, you can configure a Source transformation to represent a single Microsoft Azure SQL Data Warehouse V2 object. When you read data from Microsoft Azure SQL Data Warehouse, you can configure key range partitioning to optimize the mapping performance at run time.

The following table describes the Microsoft Azure SQL Data Warehouse V2 properties that you can configure in a Source transformation:

Property	Description
Connection	Name of the source connection.
Source Type	Source type. Select Single Object or Parameter.
Object	Name of the source object.
Sort	Sort is not supported for Microsoft Azure SQL Data Warehouse V2 Connector.

The following table describes the Microsoft Azure SQL Data Warehouse V2 source advanced properties that you can configure in a Source transformation:

Property	Description
Azure Blob Container Name	Microsoft Azure Storage container name.
Field Delimiter	Character used to separate fields in the file. Default is 0x1e. You can select the following field delimiters from the list: `~ ` ` | . TAB 0x1e`
Number of Concurrent Connections to Blob Store	Number of concurrent connections to extract data from the Microsoft Azure Blob Storage. Default is 4.
SQL Override	When you read data from a Microsoft Azure SQL Data Warehouse object, you can configure SQL overrides and define constraints.
Quote Character	The Secure Agent Skips the specified character when you read data from Microsoft Azure SQL Data Warehouse.
Pre-SQL	Pre-SQL command that must be run before reading data from the source.
Post- SQL	Post-SQL command that must be run after writing data to the target.
Tracing Level	NA

Key Range Partitioning

You can configure key range partitioning when you use a mapping task to read data from Microsoft Azure SQL Data Warehouse sources. With key range partitioning, the Secure Agent distributes rows of source data based on the fields that you define as partition keys. The Secure Agent compares the field value to the range values for each partition and sends rows to the appropriate partitions.

Use key range partitioning for columns that have an even distribution of data values. Otherwise, the partitions might have unequal size. For example, a column might have 10 rows between key values 1 and 1000 and the column might have 999 rows between key values 1001 and 2000. If the mapping includes multiple sources, use the same number of key ranges for each source.

When you define key range partitioning for a column, the Secure Agent reads the rows that are within the specified partition range. For example, if you configure two partitions for a column with the ranges as 10

through 20 and 30 through 40, the Secure Agent does not read the rows 20 through 30 because these rows are not within the specified partition range.

You can configure a partition key for fields of the following data types:

- Any type of number data type. However, you cannot use decimals in key range values.

You cannot use key range partitions when a mapping includes any of the following transformations:

- Web Services
- XML to Relational

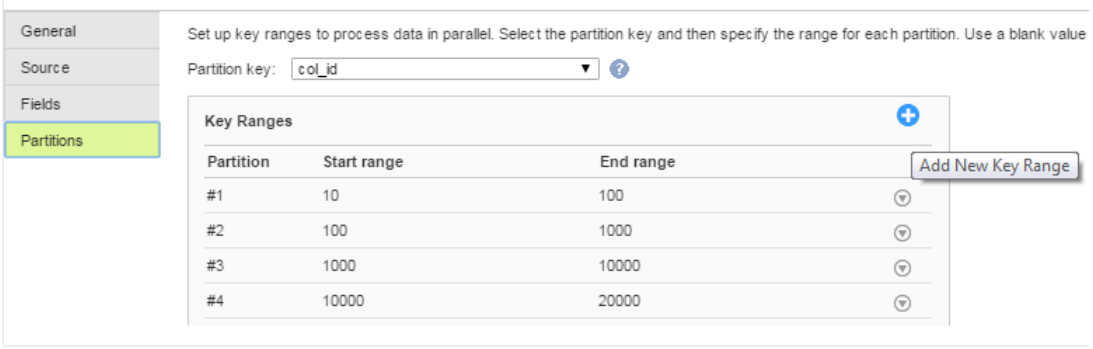
Configuring Key Range Partitioning

Perform the following steps to configure key range partitioning for Microsoft Azure SQL Data Warehouse V2 sources:

1. In the Source Properties, click the **Partitions** tab.
2. Select the required partition key from the list.
3. Click **Add New Key Range** to define the number of partitions and the key ranges based on which the Secure Agent must partition data.

   Use a blank value for the start range to indicate the minimum value. Use a blank value for the end range to indicate the maximum value.

   The following image displays the **Partitions** tab:



Microsoft Azure SQL Data Warehouse V2 Lookups in Mapping Configuration Tasks

In a mapping, you can configure a Lookup transformation to represent an Microsoft Azure SQL Data Warehouse V2 object.

When you use Microsoft Azure SQL Data Warehouse V2 object as a lookup, you need to specify the Azure Blob Container Name, Field Delimiter, and Number of Concurrent Connections to Blob Store properties.

When you use a cache lookup with Microsoft Azure SQL Data Warehouse V2 connection, the lookup condition is ignored if the lookup condition contains a NULL value.

Microsoft Azure SQL Data Warehouse V2 Target Properties in Mappings and Mapping Configuration Tasks

In a mapping, you can configure a Target transformation to represent a single Microsoft Azure SQL Data Warehouse V2 object. When the source is partitioned and you write data to Microsoft Azure SQL Data Warehouse, the Secure Agent uses the pass-through partitioning to optimize the mapping performance at run time.

The following table describes the Microsoft Azure SQL Data Warehouse V2 properties that you can configure in a Target transformation:

Property	Description
Connection	Name of the target connection.
Target Type	Target type. Select one of the following target types: - Single Object. - Parameter. Select **Parameter** to define the target type when you configure the task.
Object	Name of the target object.
Create Target	Creates a target. Enter a name for the target object and select the source fields that you want to use. By default, all source fields are used.
Operation	Select **Insert**, **Delete** , **Upsert** or **Update**.
Update Column	The key columns to upsert or update data to or delete data from Microsoft Azure SQL Data Warehouse.

The following table describes the Microsoft Azure SQL Data Warehouse V2 properties that you can configure in a Target transformation:

Advanced Property	Description
Azure Blob Container Name	Microsoft Azure Storage container name.
Field Delimiter	Character used to separate fields in the file. Default is 0x1e. You can select the following field delimiters from the list: `~ ` ` | . TAB 0x1e`
Number of Concurrent Connections to Blob Storage	Number of threads to use to move data to the staging area in Microsoft Azure Blob Storage. Default is 4.
Truncate Table	Truncates the target before inserting data to the target.
Quote Character	Specifies the quote character to skip when you write data to Microsoft Azure SQL Data Warehouse. When you write data to Microsoft Azure SQL Data Warehouse and the source table contains the specified quote character, the task fails. Change the quote character value to a value that does not exist in the source table.

Advanced Property	Description
Pre-SQL	Pre-SQL command that must be run before reading data from the source.
Post-SQL	Post-SQL command that must be run after writing data to the target.
Success File Directory	NA
Error File Directory	NA
Forward Rejected Rows	Determines whether the transformation passes rejected rows to the next transformation or drops rejected rows. By default, the Mapping Configuration application forwards rejected rows to the next transformation.

CHAPTER 5

SQL ELT Optimization

This chapter includes the following topics:

- SQL ELT Optimization Overview, 21
- Install the ODBC Driver, 21
- Create a System DSN, 22
- Create an ODBC Connection, 25
- SQL ELT Optimization Functions, 26
- Rules and Guidelines for Functions in SQL ELT Optimization, 29

SQL ELT Optimization Overview

You can use SQL ELT optimization to push transformation logic to source databases or target databases. Use SQL ELT optimization to improve task performance by using the database resources. When you run a task configured for SQL ELT optimization, the task converts the transformation logic to an SQL query. The task sends the query to the database, and the database executes the query.

Install the ODBC Driver

Before you establish an ODBC connection to connect to Microsoft Azure SQL Data Warehouse, you must install the Microsoft ODBC Driver 13 for SQL Server for Windows. Install the Microsoft ODBC Driver 13 for SQL Server on the machine where the Secure Agent is installed.

Download the Microsoft ODBC driver from the following link:
https://www.microsoft.com/en-us/download/details.aspx?id=53339

Create a System DSN

Before creating an ODBC connection, create a system DSN.

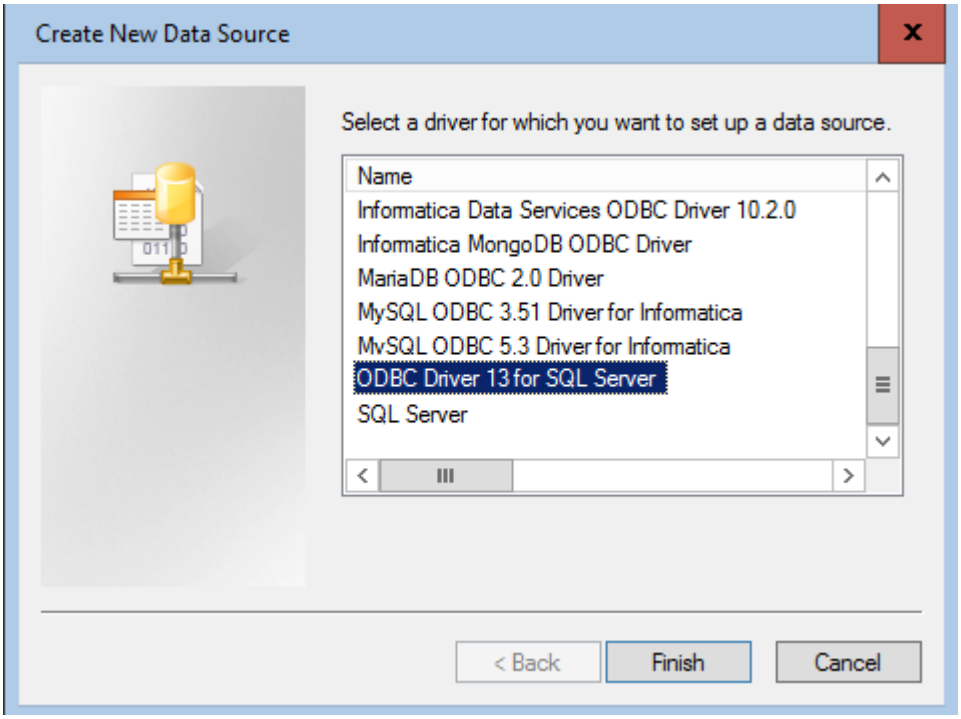
Perform the following steps to create a system DSN on the machine where the Secure Agent is installed:

1. Double-click the `odbcad32.exe` file under `C:\WINDOWS\System32`.

   The **ODBC Data Sources Administrator** box appears.
2. Click **System DSN**.

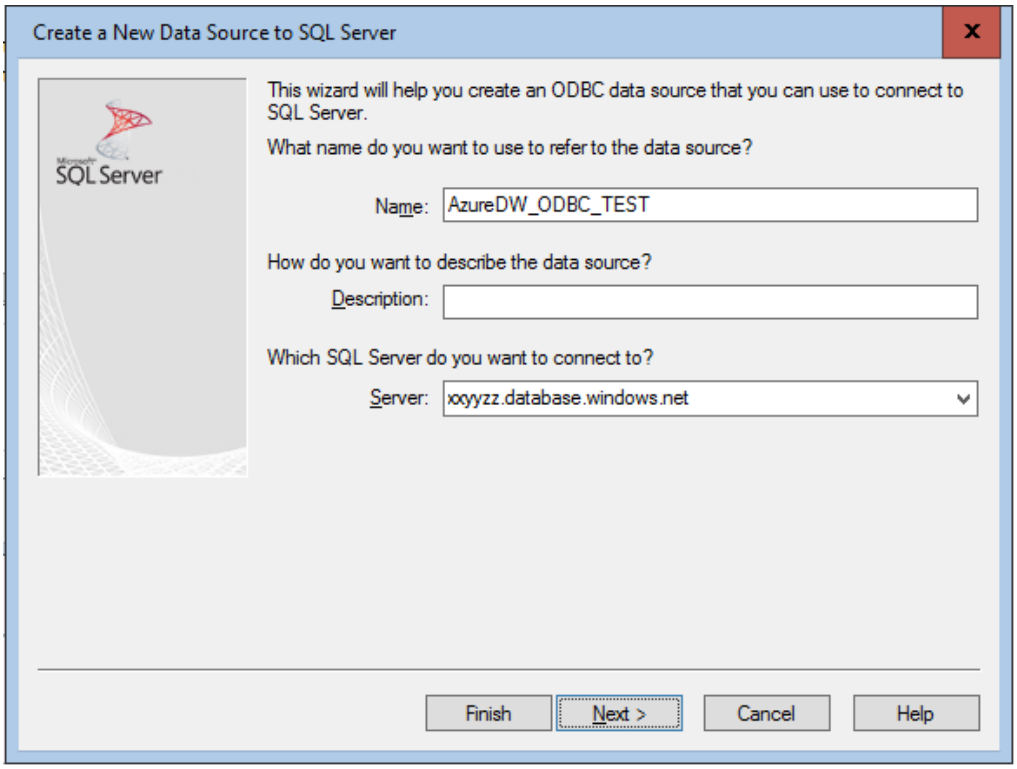
   The **System DSN** tab appears.
3. Click **Add**.

   The **Create New Data Source** box appears.



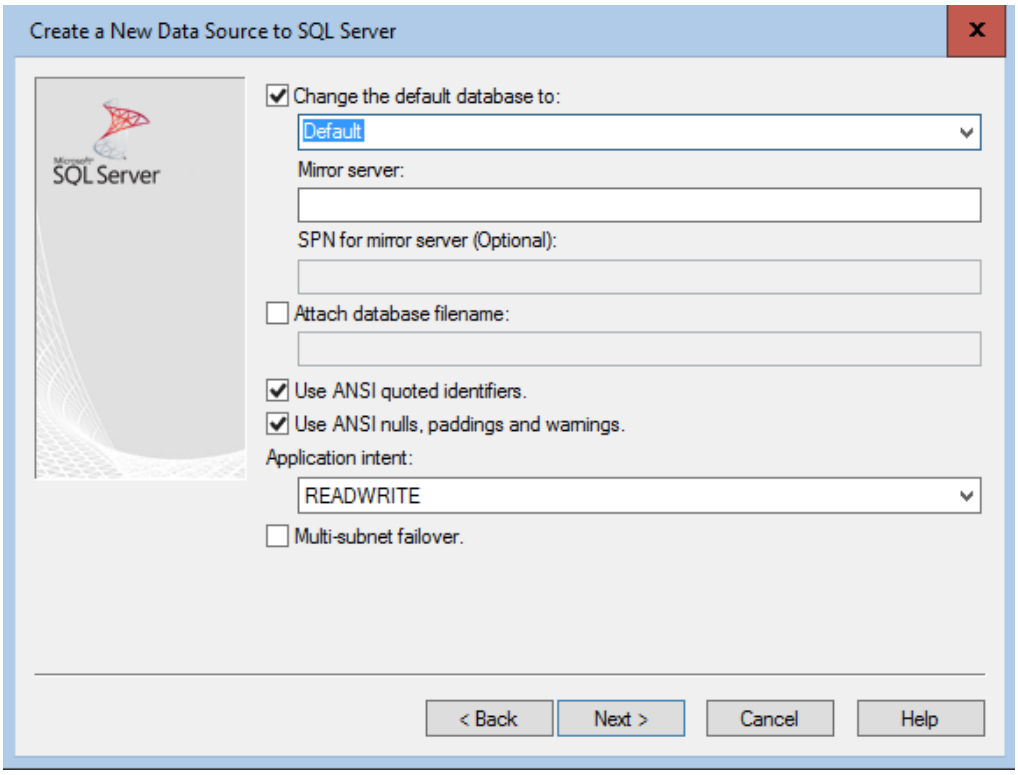
4. Select **ODBC Driver 13 for SQL Server** and click **Finish**.

The Create a New Data Source to SQL Server box appears.

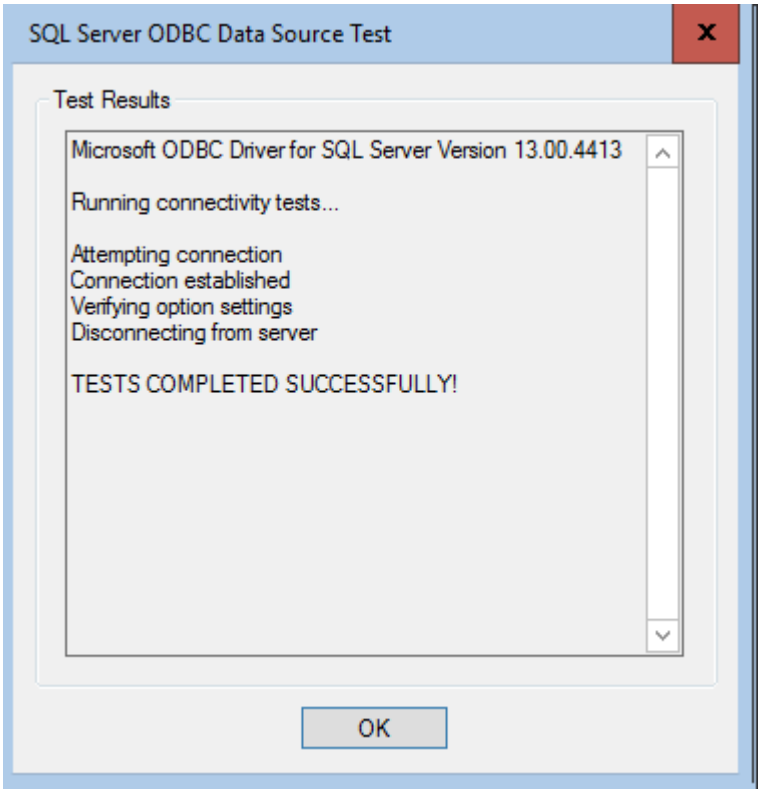


5. Specify the name, description, and the server you want to connect to.
6. Click **Next** and provide the user name and password for SQL Server Authentication.

7. Click **Next** and change the default database setting to the database you want to connect to.



8. Click **Next**.
9. Click **Finish** and then, click **Test Data Source**. Ensure that the test connection completes successfully.



10. Click **OK**.

Create an ODBC Connection

Create ODBC connections to access Microsoft Azure SQL Data Warehouse source and target objects.

Perform the following steps to create an ODBC connection on the Connections page:

1. On the Connections page, click **New Connection**. The New Connection page appears.
2. Configure the following connection details in the Connection Details section:

Property	Description
Connection Name	Name of the ODBC connection.
Description	Description of the connection.
Type	Type of the connection. Select the type of the connection as ODBC.

3. Configure the following connection details in the Connection Properties section:

Property	Description
Runtime Environment	The name of the runtime environment where you want to run the tasks.
User Name	User name of the Microsoft Azure SQL Data Warehouse account.
Password	Password for the Microsoft Azure SQL Data Warehouse account.
Data Source Name	Enter the name of the ODBC data source name that you created for the Microsoft Azure SQL Data Warehouse database.
Schema	Microsoft Azure SQL Data Warehouse schema name.
Code Page	Select the code page that the Secure Agent must use to read or write data.
ODBC Subtype	Enter the value of the ODBC Subtype field as Azure DW.

4. Click **OK**.

   The Microsoft Azure SQL Data Warehouse ODBC connection is created successfully.

Cross-Schema SQL ELT Optimization

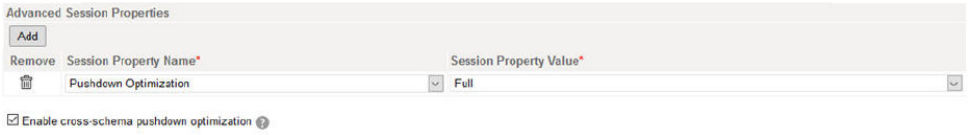
You can use cross-schema SQL ELT optimization for a mapping to read or write data based on different schemas within the same database.

To use cross-schema SQL ELT optimization, you must create two Microsoft Azure SQL Data Warehouse ODBC connections and specify different schemas for each of the connections. Select the Enable cross-schema SQL ELT optimization check box in the advanced session properties of the mapping. By default, the check box is selected.

Note: Ensure that the schemas you specify in both of the Microsoft Azure SQL Data Warehouse ODBC connections belong to the same database.

Consider the following steps to use cross-schema SQL ELT optimization:

1. Create an `Azure_DW1` ODBC connection and specify `DW_schema1` schema in the connection properties.
2. Create another `Azure_DW2` ODBC connection and specify `DW_schema2` schema in the connection properties.
3. Create a mapping and perform the following tasks:
   a. Select the `Azure_DW1` connection as source to read data from the source object using the `DW_schema1` schema.
   b. Select the `Azure_DW2` connection as target to write data to the target object using the `DW_schema2` schema.
   c. Select the Enable cross-schema SQL ELT optimization check box in the advanced session properties.
   The following image shows the Enable cross-schema SQL ELT optimization property:



   d. Run the mapping.

When you run the mapping, the Secure Agent reads data from the source object using `DW_schema1` and writes data to the target object using `DW_schema2`.

SQL ELT Optimization Functions

Microsoft Azure SQL Data Warehouse V2 supports full SQL ELT optimization

The following table summarizes the availability of SQL ELT functions in a Microsoft Azure SQL Data Warehouse. Columns marked with an X indicate that the function can be pushed to the Microsoft Azure SQL Data Warehouse by using full SQL ELT optimization. Columns marked with a dash (-) symbol indicate that the function cannot be pushed to the database.

Function	SQL ELT	Function	SQL ELT	Function	SQL ELT
ABORT()	-	INITCAP()	-	REG_MATCH()	-
ABS()	X	INSTR()	X	REG_REPLACE	-
ADD_TO_DATE()	-	IS_DATE()	-	REPLACECHR()	X
AES_DECRYPT()	-	IS_NUMBER()	-	REPLACESTR()	X
AES_ENCRYPT()	-	IS_SPACES()	-	REVERSE()	-
ASCII()	X	ISNULL()	X	ROUND(DATE)	-
AVG()	X	LAST()	X	ROUND(NUMBER)	X

Function	SQL ELT	Function	SQL ELT	Function	SQL ELT
CEIL()	X	LAST_DAY()	-	RPAD()	-
CHOOSE()	-	LEAST()	-	RTRIM()	X
CHR()	X	LENGTH()	X	SET_DATE_PART()	-
CHRCODE()	-	LN()	-	SIGN()	X
COMPRESS()	-	LOG()	-	SIN()	X
CONCAT()	-	LOOKUP	-	SINH()	-
COS()	X	LOWER()	X	SOUNDEX()	X
COSH()	-	LPAD()	-	SQRT()	X
COUNT()	X	LTRIM()	X	STDDEV()	X
CRC32()	-	MAKE_DATE_TIME()	-	SUBSTR()	-
CUME()	-	MAX()	X	SUM()	X
DATE_COMPARE()	-	MD5()	X	SYSDATE()	-
DATE_DIFF()	X	MEDIAN()	-	SYSTIMESTAMP()	X
DECODE()	X	METAPHONE()	-	TAN()	X
DECODE_BASE64()	-	MIN()	X	TANH()	-
DECOMPRESS()	-	MOD()	X	TO_BIGINT	X
ENCODE_BASE64()	-	MOVINGAVG()	-	TO_CHAR(DATE)	X
EXP()	X	MOVINGSUM()	-	TO_CHAR(NUMBER)	X
FIRST()	X	NPER()	-	TO_DATE()	X
FLOOR()	X	PERCENTILE()	-	TO_DECIMAL()	X
FV()	-	PMT()	-	TO_FLOAT()	X
GET_DATE_PART()	X	POWER()	X	TO_INTEGER()	-
GREATEST()	-	PV()	-	TRUNC(DATE)	-
IIF()	X	RAND()	-	TRUNC(NUMBER)	X
IN()	-	RATE()	-	UPPER()	X
INDEXOF()	-	REG_EXTRACT()	-	VARIANCE()	X

Microsoft Azure SQL Data Warehouse connector supports the following date formats with the `TO_DATE()` function:

- `YYYY-MM-DD HH24:MI:SS.NS`
- `YYYY-MM-DD HH12:MI:SS.NSAM`
- `MON DD YYYY HH12:MI:SS.NSAM`
- `MON DD YYYY HH24:MI:SS.NS`
- `DD MON YYYY HH12:MI:SS.NSAM`
- `DD MON YYYY HH24:MI:SS.NS`
- `MM/DD/YY HH12:MI:SS.NSAM`
- `MM/DD/YY HH24:MI:SS.NS`
- `MM/DD/YYYY HH12:MI:SS.NSAM`
- `MM/DD/YYYY HH24:MI:SS.NS`
- `HH24:MI:SS.NS`
- `HH12:MI:SS.NSAM`

Note: Use NS in the date format to support microseconds. Microsoft Azure SQL Data Warehouse supports fractional seconds upto 7 digits, which is100 nano seconds.

The following table lists the SQL ELT operators that can be used in a Microsoft Azure SQL Data Warehouse. Columns marked with an X indicate that the operator can be pushed to the Microsoft Azure SQL Data Warehouse by using full SQL ELT optimization. Columns marked with a dash (-) symbol indicate that the operator cannot be pushed to the database.

Operator	SQL ELT
+ - * /	X
%	X
||	-
= > < >= <= <>	X
!=	X
^=	X
NOT AND OR	X

The following table summarizes the availability of variables in Microsoft Azure SQL Data Warehouse. Columns marked with an X indicate that the variable can be pushed to the Microsoft Azure SQL Data

Warehouse by using full SQL ELT optimization. Columns marked with a dash (-) symbol indicate that the variable cannot be pushed to the Microsoft Azure SQL Data Warehouse.

Variable	SQL ELT
SESSSTARTTIME	-
SYSDATE	X
WORKFLOWSTARTTIME	-

Microsoft Azure SQL Data Warehouse V2 supports the following transformations for full SQL ELT optimization:

- Aggregator transformation
- Expression transformation
- Filter transformation
- Joiner transformation
- Lookup transformation
- Router transformation
- Union transformation

Rules and Guidelines for Functions in SQL ELT Optimization

Use the following rules and guidelines when pushing functions to a Microsoft Azure SQL Data Warehouse database:

- Upsert operations for SQL ELT optimization are not applicable when you use an ODBC connection.
- The `datetimeoffset` datatype is applicable only in passthrough mappings.
- The Microsoft Azure SQL Data Warehouse aggregate functions accept only one argument, which is a field set for the aggregate function. The agent ignores any filter condition defined in the argument. In addition, ensure that all fields mapped to the target are listed in the `GROUP BY` clause.
- To push the `TO_CHAR()` function to the Microsoft Azure SQL Data Warehouse database, you must define the date and format arguments.
- When you push the `SYSTIMESTAMP()` and `SYSDATE()` functions to the Microsoft Azure SQL Data Warehouse database, do not specify any format. The Microsoft Azure SQL Data Warehouse database returns the complete time stamp. `SYSDATE` works without brackets `()` only, if used it shows as invalid expression.
- You cannot push the `TO_BIGINT()` or `TO_INTEGER()` function with more than one argument to the Microsoft Azure SQL Data Warehouse database.
- When you push the `REPLACECHR()` or `REPLACESTR()` function to the Microsoft Azure SQL Data Warehouse database, the agent ignores the caseFlag argument.
  For example, both `REPLACECHR(false, in_F_CHAR, 'a', 'b')` and `REPLACECHR(true, in_F_CHAR, 'a', 'b')` return the same value.

- To push `INSTR()` to Microsoft Azure SQL Data Warehouse database, you must only define string, search_value, and start arguments. Microsoft Azure SQL Data Warehouse does not support occurrence and comparison_type arguments.
- Microsoft Azure SQL Data Warehouse connector supports the following date formats with the `TO_DATE()` function:

  - `YYYY-MM-DD HH24:MI:SS.NS`
  - `YYYY-MM-DD HH12:MI:SS.NSAM`
  - `MON DD YYYY HH12:MI:SS.NSAM`
  - `MON DD YYYY HH24:MI:SS.NS`
  - `DD MON YYYY HH12:MI:SS.NSAM`
  - `DD MON YYYY HH24:MI:SS.NS`
  - `MM/DD/YY HH12:MI:SS.NSAM`
  - `MM/DD/YY HH24:MI:SS.NS`
  - `MM/DD/YYYY HH12:MI:SS.NSAM`
  - `MM/DD/YYYY HH24:MI:SS.NS`
  - `HH24:MI:SS.NS`
  - `HH12:MI:SS.NSAM`
- To push the `SET_DATE_PART()` function to the Microsoft Azure SQL Data Warehouse database, you must use the following date data types as arguments:

  - `datetime`
  - `datetimeoffset`
  - `datetime2`
  - `smalldatetime`

  You can use the following formats for date data types:

  - `YYYY, YY, YY, Y`
  - `MM, MON, MONTH`
  - `D, DD, DD, DY, DAY`
  - `HH, HH12, HH24`
  - `MI`
  - `MS`
  - `SS`

  Note: `NS` and `US` formats are not applicable to `SET_DATE_PART()`.
- To push the `ADD_TO_DATE()` function to the Microsoft Azure SQL Data Warehouse database, you must use the following date data types as arguments:

  - `date`
  - `datetime`
  - `datetimeoffset`
  - `datetime2`
  - `smalldatetime`
  - `time`

You can use the following formats for date data types:

- `YYYY, YY, YY, Y`
- `MM, MON, MONTH`
- `D, DD, DD, DY, DAY`
- `HH, HH12, HH24`
- `MI`
- `MS`
- `SS`
- `NS`: applicable to `datetimeoffset`, `datetime2`, and `time`
- `US`

- To push the `MAKE_DATE_TIME()` function to the Microsoft Azure SQL Data Warehouse database, you must use the following date data types as arguments:

  - `date`
  - `datetime`
  - `datetimeoffset`
  - `datetime2`
  - `smalldatetime`
  - `time`

  You can use year, month, day, hour, minute, second, and nanosecond with appropriate return date types.

CHAPTER 6

Data Type Reference

This chapter includes the following topics:

- Microsoft Azure SQL Data Warehouse V2 Data Type Reference Overview, 32
- Microsoft Azure SQL Data Warehouse V2 and Transformation Data Types, 33

Microsoft Azure SQL Data Warehouse V2 Data Type Reference Overview

Microsoft Azure SQL Data Warehouse V2 native data types

Microsoft Azure SQL Data Warehouse V2 data types appear in the Fields tab for Source and Target transformations when you choose to edit metadata for the fields.

Transformation data types

Set of data types that appear in the remaining transformations. They are internal data types based on ANSI SQL-92 generic data types, which Data Integration uses to move data across platforms. Transformation data types appear in all remaining transformations in mappings, synchronization tasks, or mapping tasks.

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

Microsoft Azure SQL Data Warehouse V2 and Transformation Data Types

The following table compares the Microsoft Azure SQL Data Warehouse V2 native data type to the transformation data type:

Microsoft SQL Data Warehouse Data Type	Transformation Data Type
bigint	bigint
binary	binary
bit	integer
char	string
date	date/time
datetime	date/time
datetime2	date/time
decimal	decimal
float	decimal
int	integer
money	double
nchar	string
nvarchar	string
real	decimal
smalldatetime	date/time
smallint	integer
smallmoney	double
time	String
tinyint	integer
varbinary	binary
varchar	string

INDEX

A

administration 8
advanced source properties 12
advanced target properties 14
Azure active directory authentication 8
Azure Blob Container Name 16, 19

C

Cloud Application Integration community
  URL 5
Cloud Developer community
  URL 5
connections
  Microsoft Azure SQL Data Warehouse V2 11
create
  system DSN 22

D

Data Integration community
  URL 5
data types 32, 33
delete 7

E

error file directory 14, 19

F

field delimiter 16, 19
forward rejected rows 19

I

IDENTITY column 12
Informatica Global Customer Support
  contact information 6
Informatica Intelligent Cloud Services
  web site 5
insert 7
install
  Microsoft ODBC driver 21

L

lookups 8

M

maintenance outages 6
mapping tasks 8
Mapping tasks
  overview 16
mappings
  overview 16
  source properties 16
  target properties 19
Microsoft Azure SQL Data Warehouse V2
  connection properties 11

N

native data type 33
number of concurrent connections to Blob Storage 12
number of concurrent connections to Blob Store 14
number of concurrent sessions to Blob Storage 16, 19

O

operations 7

P

properties
  in mappings 16, 19

R

runtime environments 8

S

Secure Agents 8
sources 8
status
  Informatica Intelligent Cloud Services 6
success file directory 14, 19
synchronization tasks 8
Synchronization tasks
  overview 12
system status 6

T

targets 8
tracing level 16
transformation data type 33

trust site
description 6

U

upgrade notifications 6

W

web site 5