Informatica®

Informatica® Cloud Data Integration

Microsoft Fabric Lakehouse Connector

Informatica Cloud Data Integration Microsoft Fabric Lakehouse Connector
July 2024

© Copyright Informatica LLC 2023, 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: 2024-07-15

Table of Contents

Preface 4

Informatica Resources 4

- Informatica Documentation 4
- Informatica Intelligent Cloud Services web site 4
- Informatica Intelligent Cloud Services Communities 4
- Informatica Intelligent Cloud Services Marketplace 4
- Data Integration connector documentation 5
- Informatica Knowledge Base 5
- Informatica Intelligent Cloud Services Trust Center 5
- Informatica Global Customer Support 5

Chapter 1: Introduction to Microsoft Fabric Lakehouse Connector 6

Microsoft Fabric Lakehouse Connector assets 6

Administration of Microsoft Fabric Lakehouse Connector 6

Chapter 2: Microsoft Fabric Lakehouse connections 8

Microsoft Fabric Lakehouse connection properties 8

Chapter 3: Mappings and mapping tasks with Microsoft Fabric Lakehouse Connector 10

Microsoft Fabric Lakehouse sources in mappings 11

Microsoft Fabric Lakehouse targets in mappings 12

Rules and guidelines for mappings 13

Chapter 4: Migrating a mapping 16

Migrate a mapping within the same path 16

Migrate a mapping to a different path 16

- Migration options 17
- Migration restrictions 18

Chapter 5: Data type reference 19

Microsoft Fabric Lakehouse and transformation data types 19

Index 21

Preface

Use Microsoft Fabric Lakehouse Connector to learn how to read or write to Microsoft Fabric Lakehouse by using Data Integration. Learn to create a Microsoft Fabric Lakehouse connection, develop, and run mappings and mappings in advanced mode, and tasks in Data Integration.

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 Microsoft Fabric Lakehouse Connector

You can use Microsoft Fabric Lakehouse Connector to securely read data from and write data to Microsoft Fabric Lakehouse from Data Integration.

Create a Microsoft Fabric Lakehouse connection and use the connection as a source and as a target in mappings and mapping tasks.

You can switch mappings to advanced mode to include transformations and functions that enable advanced functionality. A mapping in advanced mode can run on an advanced cluster hosted on Microsoft Azure or a self-service cluster.

Microsoft Fabric Lakehouse Connector assets

Create assets in Data Integration to integrate data using Microsoft Fabric Lakehouse Connector.

When you use Microsoft Fabric Lakehouse Connector, you can include the following Data Integration assets:

- Mapping
- Mapping task

Administration of Microsoft Fabric Lakehouse Connector

Before you use Microsoft Fabric Lakehouse Connector, an administrator needs to perform certain prerequisite tasks.

1. Configure the following prerequisites in Power BI:
   a. Create a workspace.
   b. Create a lakehouse in the workspace. Ensure that you have admin privileges to create the lakehouse.
   For more information about how to create a workspace and lakehouse, see the Microsoft Power BI documentation.

2. Configure the following prerequisites in the Azure portal:
    a. Create an application in Azure Active Directory and generate the client ID, tenant ID, and client secret.
    b. Create a security group in Azure Active Directory and add owners or members to the security group. For more information about how to add owners or members to the security group in Azure Active Directory, see the Microsoft Azure documentation.
3. Enable the Power BI service admin settings to allow the Azure Active Directory application to use the Power BI APIs.
4. Add the Azure Active Directory application to the workspace in Power BI.

CHAPTER 2

Microsoft Fabric Lakehouse connections

Create a Microsoft Fabric Lakehouse connection to securely read data from or write data to Microsoft Fabric Lakehouse.

You can use a Microsoft Fabric Lakehouse connection to specify sources and targets in mappings and mapping tasks.

Microsoft Fabric Lakehouse connection properties

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

The following table describes the basic connection properties:

Property	Description
Connection Name	Name of the connection. Each connection name must be unique within the organization. Connection names can contain alphanumeric characters, spaces, and the following special characters: _ . + -, Maximum length is 255 characters.
Description	Description of the connection. Maximum length is 4000 characters.
Type	Microsoft Fabric Lakehouse
Runtime Environment	The name of the runtime environment where you want to run tasks. Select a Secure Agent, Hosted Agent, or serverless runtime environment.
SQL Connection String	The SQL connection string to connect to Microsoft Fabric Lakehouse. Specify the connection string in the following format: `<Server>.lakehouse.pbidedicated.windows.net`
Client ID	The application ID or client ID of your application registered in Azure Active Directory for service principal authentication.
Client Secret	The client secret for your application registered in Azure Active Directory.
Tenant ID	The tenant ID of your application registered in Azure Active Directory.

Property	Description
Workspace	The name of the workspace in Microsoft Fabric Lakehouse that you want to connect.
Database	The name of the database in Microsoft Fabric Lakehouse that you want to connect.

CHAPTER 3

Mappings and mapping tasks with Microsoft Fabric Lakehouse Connector

You can configure a mapping to read from and write to Microsoft Fabric Lakehouse tables.

You can configure a Source or Target transformation to represent a Microsoft Fabric Lakehouse object. Add the Source and Target transformations in the mapping canvas and configure the Microsoft Fabric Lakehouse source and target properties. Use the Data Integration Mapping Designer to define and configure a mapping for Microsoft Fabric Lakehouse sources and targets.

In advanced mode, the Mapping Designer updates the mapping canvas to include transformations and functions that enable advanced functionality. When you run a mapping in advanced mode, Data Integration analyzes the data logic defined in the mapping to automatically assign data logic to run on an elastic cluster.

You can use Monitor to monitor the job.

Microsoft Fabric Lakehouse sources in mappings

To read data from Microsoft Fabric Lakehouse, configure a Microsoft Fabric Lakehouse object as a source in a mapping or mapping in advanced mode.

The following table describes the source properties that you can configure for a Microsoft Fabric Lakehouse source:

Property	Description
Connection	Name of the source connection. Select an existing connection, create a new connection, or define parameter values for the source connection property. If you want to overwrite the parameter at runtime, select the **Allow parameter to be overridden at run time** option when you create a parameter. When the task runs, the agent uses the parameters from the file that you specify in the task advanced session properties. Ensure that the parameter file is in the correct format. Specify the parameter file directory and name in the advanced session properties. When you switch between a non-parameterized and a parameterized Microsoft Fabric Lakehouse connection, the advanced property values are retained.
Source Type	Type of the source object. Select Single Object, Query, or Parameter. Default is Single Object.
Parameter	A parameter file where you define values that you want to update without having to edit the task. The **Parameter** property appears only if you select the source type as **Parameter**. Select an existing parameter for the source object or click **New Parameter** to define a new parameter for the source object. If you want to overwrite the parameter at runtime, select the **Allow parameter to be overridden at run time** option when you create a parameter. When the task runs, the agent uses the parameters from the file that you specify in the task advanced session properties. Ensure that the parameter file is in the correct format.
Object	Name of the Microsoft Fabric Lakehouse source object. Select the source object for a single source. You can also select an existing object from the list. Ensure that the table headers of the source object do not contain special characters.

The following table describes the query options that you can configure for a Microsoft Fabric Lakehouse source:

Property	Description
Filter	Filters record based on the filter condition in a read operation. Click **Configure** to add conditions to filter records and reduce the number of rows that the Secure Agent reads from the source. Select one of the following filter conditions from the list: - Not parameterized. Uses a basic filter to specify the object, field, operator, and value to select specific records. - Completely parameterized. Not applicable. - Advanced. Not applicable. Default is Not parameterized.
Sort	Not applicable.

The following table describes the advanced source properties that you can configure for a Microsoft Fabric Lakehouse source:

Property	Description
Table Name Override	Overrides the table name configured in the source object and uses the specified table name to read data at runtime.
Tracing Level	Sets the amount of detail that appears in the log file. Select one of the following options from the list: - Terse - Normal - Verbose Initialization - Verbose Data Default is Normal.

Microsoft Fabric Lakehouse targets in mappings

To write data to a Microsoft Fabric Lakehouse, configure a Microsoft Fabric Lakehouse object as a target in a mapping or mapping in advanced mode.

The following table describes the target properties that you can configure for a Microsoft Fabric Lakehouse target:

Property	Description
Connection	Name of the target connection. Select an existing connection, create a new connection, or define parameter values for the target connection property. If you want to overwrite the target connection properties at runtime, select the **Allow parameter to be overridden at run time** option when you create a parameter. When the task runs, the agent uses the parameters from the file that you specify in the task advanced session properties. Ensure that the parameter file is in the correct format. Specify the parameter file directory and name in the advanced session properties. When you switch between a non-parameterized and a parameterized Microsoft Fabric Lakehouse connection, the advanced property values are retained.
Target Type	Type of the Microsoft Fabric Lakehouse target object. Select one of the following source types from the list: - Single Object. Select to specify a single Microsoft Fabric Lakehouse object. - Parameter. Select to specify a parameter name. You can configure the target object in a mapping task associated with a mapping or mapping in advanced mode that uses this Target transformation. Default is Single Object.

Property	Description
Parameter	A parameter file where you define values that you want to update without having to edit the task. The **Parameter** property appears only if you select the target type as **Parameter**. Select an existing parameter for the target object or click **New Parameter** to define a new parameter for the target object. If you want to overwrite the target connection properties at runtime, select the **Allow parameter to be overridden at run time** option when you create a parameter. When the task runs, the agent uses the parameters from the file that you specify in the task advanced session properties. Ensure that the parameter file is in the correct format.
Object	Name of the Microsoft Fabric Lakehouse target object. Specify the target object that you want to create. You can also select an existing object from the list.
Operation	Type of target operation. Select Insert, Update, or Delete. Default is Insert. You can only perform an insert operation in a mapping. You cannot perform data driven and upsert target operations in a mapping or mapping in advanced mode.

The following table describes the advanced target properties that you can configure for a Microsoft Fabric Lakehouse target:

Property	Description
Update Mode	Loads data to the target based on the mode you select. This property is applicable when you select the **Update** target operation. Select one of the following modes: - Update As Update. Updates all rows flagged for update if the entries exist. - Update Else Insert. The agent first updates all rows flagged for update if the entries exist in the target. If the entries do not exist, the agent inserts the entries.
Table Name Override	Overrides the table name configured in the target object and uses the specified table name to write data at runtime.
Forward Rejected Rows	Determines whether the transformation passes rejected rows to the next transformation or drops rejected rows. By default, the mapping task forwards rejected rows to the next transformation. If you select the Forward Rejected Rows option, the Secure Agent flags the rows for reject and writes them to the reject file. If you do not select the Forward Rejected Rows option, the Secure Agent drops the rejected rows and writes them to the session log file. The Secure Agent does not write the rejected rows to the reject file.

Rules and guidelines for mappings

Consider the following rules and guidelines for Microsoft Fabric Lakehouse mappings:

General guidelines

- When you read from Microsoft Fabric Data Warehouse and write to Microsoft Fabric Lakehouse, the mapping fails.
- You can't configure partitioning in a mapping.
- When you specify a custom query for a source object, ensure that the query does not contain a semicolon at the end of the SQL statement. Otherwise, the mapping fails.
- You cannot use input parameters in custom queries that read data from Microsoft Fabric Lakehouse source objects.
- When you use a stored procedure in a custom query to read data from Microsoft Fabric Lakehouse source objects, the mapping fails at runtime.
- Ensure that you create the parameter before you override it using a parameter file.
- You can't parameterize a target created at runtime using the parameter file.
- A mapping that uses a MINUS operator in a custom query to read data from Microsoft Fabric Lakehouse source objects fails at runtime. Use an EXCEPT operator instead of a MINUS operator.
- When you use an ORDER BY clause along with Join or Union clauses in a custom query to read data from Microsoft Fabric Lakehouse source objects, ensure that you add "offset 0 rows" at the end of the query. Otherwise, the mapping fails.
- When you specify a table name in a custom query to read data from Microsoft Fabric Lakehouse source objects, specify the exact table name as in the source table because the table name is case sensitive.
- To avoid a Java heap size error when you write to Microsoft Fabric Lakehouse, you need to allocate more memory to the DTM in the Secure Agent properties based on the amount of data that you want to write. Increase the -Xms and -Xmx values for the DTM in the JVM options in the system configuration details of the Secure Agent. The recommended -Xms value is 512 MB, while the recommended -Xmx value is 1024 MB.
- When you filter records in a read operation, you can use only the following filter operators:
  - Equals
  - Greater
  - Greater_or_equals
  - Less_or_equals
  - Less
  - Not_Equals

Data types

- When you write data to a Microsoft Fabric Lakehouse target created at runtime, Float or Real data types are mapped to and written as Double data type in the target. To write Float or Real data types, use the **Edit Metadata** option to edit the data type in the Target transformation.
- When you read data of the Float or Real data type from a source and write to a target, the values written to the target are not accurate.
- A mapping fails in the following cases:
  - Data is of the Date data type and the date is less than 1582-10-15.
  - Data is of the Int96 data type and the timestamp is less than 1900-01-01T00:00:00Z.

To resolve this issue, specify the following spark session properties in the mapping task or in the custom properties file for the Secure Agent:

- `spark.sql.legacy.timeParserPolicy=LEGACY`
- `spark.sql.parquet.int96RebaseModeInWrite=LEGACY`
- `spark.sql.parquet.datetimeRebaseModeInWrite=LEGACY`
- `spark.sql.parquet.int96RebaseModeInRead=LEGACY`
- `spark.sql.parquet.datetimeRebaseModeInRead=LEGACY`

CHAPTER 4

Migrating a mapping

You can configure a connection and mapping in one environment and then migrate and run the mapping in another environment.

After the migration, you can change the connection properties from the Administrator service, but you do not need to modify the mapping. Data Integration uses the configured runtime attributes from the earlier environment to run the mapping successfully in the new environment.

Migrate a mapping within the same path

If you want the migrated mapping to use the same object path as in the earlier environment, you must maintain the same schema and table in the Microsoft Fabric Lakehouse account for both the earlier and the new environment.

For example, if you have two different accounts, Account1 used for Organization 1 and Account2 used for Organization 2, the object path for the schema and table must be the same in both the accounts:

Account1: SCHEMA1/TABLE1

Account2: SCHEMA1/TABLE1

You do not need to override the schema and table in the advanced properties.

Migrate a mapping to a different path

You can use a different object path to run the mapping from the new environment.

Before you migrate the mapping, you can change the object metadata, runtime attributes, or the connection attributes to reflect the object path in the migrated environment. You do not have to edit or update the mapping in the new environment.

As a rule, when you specify the schema or table in the advanced properties, connection, or object properties, Data Integration honors the attributes in the following order of precedence:

1. **Runtime advanced attributes**. The advanced properties such as schema and table in the Source or Target transformation in a mapping.
2. **Connection attributes**. Attributes such as database and schema set in the connection properties.
3. **Object metadata**. The object selected in the Source or Target transformation in a mapping.

Migration options

When you migrate to a different path, you can choose from one of the following options to update the object path:

Option 1. Update the connection properties to reference the new object

When you import the mapping from, for example, Org 1, into a new organization, for example, Org 2, you can change the existing connection in the **Review Connections** section to map to the connection that has access to the specified schema and table in Org 2.

Option 2. Override the properties from the advanced properties

Before the migration, specify the required schema and table name for the object in the new organization in the advanced properties of the Org 1 mapping.

After the migration, when you run the mapping, the Secure Agent uses the configured advanced parameters to override the object specified in the mapping imported from Org 1.

Option 3. Parameterize the properties in the mapping

You can choose to parameterize the advanced attributes, such as the schema and table name before the migration. You can configure input parameters, in-out parameters, and parameter files in the mapping. When you use a parameter file, you can save the parameter file on a local machine or in a cloud-hosted directory. After you migrate the mapping, do not edit or update the mapping. If you have used in-out parameters for the advanced attributes such as for the schema and table name, you can update these from the parameter file.

Parameterizing only the advanced properties, but not the object

You can choose to parameterize only the advanced properties before the migration, and use them at runtime. After the migration, when you run the mapping, the Secure Agent uses the value of the parameters specified in the advanced properties.

Parameterizing both the advanced properties and the object or connection

When you parameterize both the advanced properties and the object, you can choose to select or deselect the Allow parameter to be overridden at runtime option.

Allow parameter to be overridden at runtime option is selected

When you parameterize both the advanced properties and the object or connection and select the **Allow parameter to be overridden at runtime** option before the migration, you must select a placeholder object in the object properties in the mapping and then specify an override to this placeholder object from the advanced properties. Ensure that the placeholder object contains the same metadata as the corresponding table that you specify as an override. When you run the mapping, the value specified in the advanced property overrides the placeholder object.

Allow parameter to be overridden at runtime option is not selected

When you parameterize both the advanced properties and the object or connection and do not select the **Allow parameter to be overridden at runtime** option before the migration, you do not need a placeholder object and the values specified in the advanced properties take precedence.

Migration restrictions

Consider some rules for some of the configurations before you migrate assets:

- Retain the following transformation, object type, and fields in the advanced properties of a mapping:

Transformations	Object type
Source	Single object, query
Target	Single object
Note: Advanced fields such as the schema and table name are supported for the Source and Target transformations.	

- If you use in-out parameters for the table name and override the parameter from a parameter file, ensure that you define default values for the in-out parameter.

CHAPTER 5

Data type reference

Data Integration uses the following data types in mappings and mapping tasks with Microsoft Fabric Lakehouse:

Microsoft Fabric Lakehouse native data types

Microsoft Fabric Lakehouse 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 or mappings in advanced mode.

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.

Microsoft Fabric Lakehouse and transformation data types

Use only the supported Microsoft Fabric Lakehouse data types in a mapping or mapping in advanced mode to read from and write to Microsoft Fabric Lakehouse.

The following table lists the Microsoft Fabric Lakehouse data types that Data Integration supports and the corresponding transformation data types:

Microsoft Fabric Lakehouse Data Type	Transformation Data Type	Description
Bigint	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807
Bit	Integer	1, 0, or NULL
Char	String	1 to 8000 characters
Date	Date/Time	January 1, 1900 to December 31, 9999
Datetime2	Date/Time	Date range: January 1, 1900 to December 31, 9999 Time range : 00:00:00 to 23:59:59.997

Microsoft Fabric Lakehouse Data Type	Transformation Data Type	Description
Decimal	Decimal	Maximum value: 9.999999999999998e+37 Minimum value: -9.999999999999998e+37
Float	Double	Precision 15, scale 0
Int	Integer	-2,147,483,648 to 2,147,483,647
Real	Double	Precision 7, scale 0
Smallint	Integer	-32,768 to 32,767
Varbinary	Binary	1 to 8000 bytes
Varchar	String	1 to 8000 characters

INDEX

A

administration 6

C

Cloud Application Integration community
  URL 4
Cloud Developer community
  URL 4
connections 8

D

Data Integration community
  URL 4
data type reference
  overview 19

I

Informatica Global Customer Support
  contact information 5
Informatica Intelligent Cloud Services
  web site 4

M

maintenance outages 5
mapping
  Microsoft Fabric Lakehouse sources 11
  Source transformation 11
Microsoft Fabric Lakehouse
  assets 6
Microsoft Fabric Lakehouse *(continued)*
  connection properties 8
  data types 19
Microsoft Fabric Lakehouse connection
  overview 8
Microsoft Fabric Lakehouse Connector
  mapping 10
  mapping task 10
  overview 6
  rules and guidelines 13
Microsoft Fabric Lakehouse targets
  mappings 12

S

status
  Informatica Intelligent Cloud Services 5
system status 5

T

Target transformation 12
trust site
  description 5

U

upgrade notifications 5

W

web site 4