



Informatica® PowerExchange for Snowflake
10.2 HotFix 1

User Guide

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.

Informatica, the Informatica logo, PowerCenter, and PowerExchange 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.

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.

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

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

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

Table of Contents

Preface	5
Informatica Resources.	5
Informatica Network.	5
Informatica Knowledge Base.	5
Informatica Documentation.	5
Informatica Product Availability Matrixes.	6
Informatica Velocity.	6
Informatica Marketplace.	6
Informatica Global Customer Support.	6
 Chapter 1: Introduction to PowerExchange for Snowflake.....	 7
PowerExchange for Snowflake Overview.	7
Introduction to Snowflake.	7
 Chapter 2: Installation and Configuration.....	 9
Installation and Configuration Overview.	9
Installing the Server Component on UNIX.	9
Installing the Client Component.	10
 Chapter 3: Snowflake Connections.....	 12
Snowflake Connections Overview.	12
Snowflake Connection Properties.	12
Creating a Snowflake Connection.	13
 Chapter 4: PowerExchange for Snowflake Data Objects.....	 14
Snowflake Data Object Overview.	14
Snowflake Data Object Properties.	14
Snowflake Data Object Read Operation.	15
Snowflake Data Object Read Operation Properties.	15
Snowflake Data Object Write Operation.	16
Snowflake Data Object Write Operation Properties.	17
Creating a Snowflake Data Object.	18
Creating a Snowflake Data Object Operation.	18
Rules and Guidelines for Snowflake Mappings.	18
 Chapter 5: PowerExchange for Snowflake Mappings.....	 20
PowerExchange for Snowflake Mappings Overview.	20
Mapping Validation and Run-time Environments.	20
Snowflake Mapping Example.	21

Appendix A: Snowflake Data Type Reference..... 22
Data Type Reference Overview. 22
Snowflake and Transformation Data Types. 22

Index..... 24

Preface

The *Informatica PowerExchange for Snowflake User Guide* provides information about extracting data from and loading data to Snowflake. This guide is written for database administrators and developers who are responsible for developing mappings that extract data from and load data to Snowflake. This guide assumes you have knowledge of Snowflake and Informatica Data Services.

Informatica Resources

Informatica Network

Informatica Network hosts Informatica Global Customer Support, the Informatica Knowledge Base, and other product resources. To access Informatica Network, visit <https://network.informatica.com>.

As a member, you can:

- Access all of your Informatica resources in one place.
- Search the Knowledge Base for product resources, including documentation, FAQs, and best practices.
- View product availability information.
- Review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to search Informatica Network for product resources such as documentation, how-to articles, best practices, and PAMs.

To access the Knowledge Base, visit <https://kb.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 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 Product Availability Matrixes

Product Availability Matrixes (PAMs) indicate the versions of operating systems, databases, and other types of data sources and targets that a product release supports. If you are an Informatica Network member, you can access PAMs at

<https://network.informatica.com/community/informatica-network/product-availability-matrices>.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services. Developed from the real-world experience of hundreds of data management projects, Informatica Velocity represents the collective knowledge of our consultants who have worked with organizations from around the world to plan, develop, deploy, and maintain successful data management solutions.

If you are an Informatica Network member, you can access Informatica Velocity resources at <http://velocity.informatica.com>.

If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that augment, extend, or enhance your Informatica implementations. By leveraging any of the hundreds of solutions from Informatica developers and partners, you can improve your productivity and speed up time to implementation on your projects. You can access Informatica Marketplace at <https://marketplace.informatica.com>.

Informatica Global Customer Support

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

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:

<http://www.informatica.com/us/services-and-training/support-services/global-support-centers>.

If you are an Informatica Network member, you can use Online Support at <http://network.informatica.com>.

CHAPTER 1

Introduction to PowerExchange for Snowflake

This chapter includes the following topics:

- [PowerExchange for Snowflake Overview, 7](#)
- [Introduction to Snowflake, 7](#)

PowerExchange for Snowflake Overview

You can use PowerExchange for Snowflake to extract data from and load data to Snowflake. You can extract data from other applications, databases, and flat files, and use PowerExchange for Snowflake to load the data to Snowflake.

You can use Snowflake objects as sources and targets in mappings. When you use Snowflake objects in mappings, you must configure properties specific to Snowflake.

You can run Snowflake mappings in the native or Hadoop environment. If you choose the Hadoop environment, you can run Snowflake mappings on the Spark engine.

Note: PowerExchange for Snowflake supports Cloudera 5.12 Kerberos, Cloudera 5.12 Non-Kerberos, and Amazon EMR 5.4 Non-Kerberos Hadoop distributions.

Example

An enterprise application uses an Oracle database to store the product transaction details such as transactionID, customerID, productID, quantity, and order date. You need to analyze the completed transactions, pending transactions, and availability of stock. Use PowerExchange for Snowflake to create a mapping to extract all the transaction records from the Oracle source, and load the records to a Snowflake target for data analysis.

Introduction to Snowflake

Snowflake is a cloud data warehouse service that organizations can use to store and analyze data.

Snowflake is a Software-as-a-Service (SaaS) application that uses an SQL database engine with an architecture designed for the cloud. Apart from the features of other enterprise data warehouses, Snowflake has additional functionalities and capabilities.

The Snowflake architecture comprises three layers:

- **Database Storage.** When you load data to Snowflake, Snowflake reorganizes and stores the data in the Snowflake database. You can access the data stored in the Snowflake database through SQL query operations that you run by using Snowflake.
- **Query Processing.** Snowflake processes all queries in the query processing layer. The processing layer contains all the compute resources that Snowflake needs to execute queries. For example, Snowflake uses CPU, memory, and temporary storage to execute queries.
- **Cloud Services.** The cloud services layer contains all the different components that Snowflake needs to process the user requests. For example, Snowflake uses authentication, infrastructure management, and access control services to process user requests.

CHAPTER 2

Installation and Configuration

This chapter includes the following topic:

- [Installation and Configuration Overview, 9](#)

Installation and Configuration Overview

The PowerExchange for Snowflake installation consists of a server installation and a client installation. You can install PowerExchange for Snowflake on UNIX.

To use PowerExchange for Snowflake, perform the following steps:

1. Install or upgrade Informatica services. Create and configure a Model Repository Service and a Data Integration Service.
2. Verify that you have write permissions on all the directories within the `<INFA_HOME>` directory.
3. Ensure that the PowerExchange for Snowflake license is activated.
4. Install the PowerExchange for Snowflake server component after you install the Informatica services. The server binaries are copied to the Informatica installation directory.
5. Install the PowerExchange for Snowflake client component after you install the Informatica clients. The client binaries are copied to the Informatica installation directory.

For more information about product requirements and supported platforms, see the Product Availability Matrix on Informatica Network:

<https://network.informatica.com/community/informatica-network/product-availability-matrices>

Installing the Server Component on UNIX

If multiple nodes exist in your environment, you must first install the server component on the master gateway node. You can then install the server component on the other nodes in the domain.

Before you install, shut down the Informatica domain.

1. Delete the contents from the following directories:
 - `$INFA_HOME/services/work_dir`
 - `$INFA_HOME/tomcat/bin/workspace`
2. Navigate to the root directory of the extracted installer files.
3. Enter `./install.sh` at the command prompt.

Note: The `install.sh` file must have executable permissions.

4. Enter the path to the Informatica installation directory.

By default, the server components are installed in the following location:

```
<User Home Directory>/Informatica/<version folder>
```

If you did not shut down the domain, a message appears asking you to shut down the domain.

5. Review the installation information and press **Enter** to begin the installation.
6. View or enter the domain information.

Property	Description
Domain Name	Name of the domain where Informatica services are installed. This field is read-only.
Node Name	Name of the node on which you are installing the PowerExchange for Snowflake server component. This field is read-only.
Domain User Name	User name of the administrator for the domain.
Domain Password	Password for the domain administrator.
Master Gateway Node	Indicates whether the node on which you are installing the server component is the master gateway node. Select from the following options: 1. Yes. Select Yes if the node is the master gateway node. 2. No. Select No for all other nodes on which you install the server component.

For more information about the tasks performed by the installer, view the installation log files.

After you complete the installation, ensure that the binary files in the `$INFA_HOME/services/shared/bin` directory have executable permissions.

Installing the Client Component

Install the client component on every Informatica Developer client machine that connects to the domain.

1. Delete the contents from the following directory:

```
$INFA_HOME\clients\DeveloperClient\workspace
```

2. Delete the configuration files and retain the `config.ini` file from the following directory:

```
$INFA_HOME\clients\DeveloperClient\configuration
```

3. Unzip the client installation archive and navigate to the root directory of the extracted installer files.

4. Run the `install.bat` script file.

The **Welcome** page appears.

5. Click **Next**.

The **Installation Directory** page appears.

6. Enter the absolute path to the Informatica installation directory. Click **Browse** to find the directory or use the default directory.

7. Click **Next**.

The **Pre-Installation Summary** page appears.

8. Verify that all installation requirements are met and click **Install**.

The installer shows the progress of the installation. When the installation is complete, the **Post-Installation Summary** page displays the status of the installation.

9. Click **Done** to close the installer.

For more information about the tasks performed by the installer, view the installation log files.

CHAPTER 3

Snowflake Connections

This chapter includes the following topics:

- [Snowflake Connections Overview, 12](#)
- [Snowflake Connection Properties, 12](#)
- [Creating a Snowflake Connection, 13](#)

Snowflake Connections Overview

Use a Snowflake connection to access a Snowflake database.

Use the Snowflake connection to import Snowflake metadata, create data objects, preview data, and run mappings. When you create a Snowflake connection, you define the connection attributes that the Developer tool uses to connect to the Snowflake database.

Use the Developer tool, Administrator tool, or infacmd to create a Snowflake connection.

Snowflake Connection Properties

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

The following table describes the Snowflake connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~`!\$%^&*()-+={} \:;"'<,>./
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.

Property	Description
Type	The connection type. Select SnowflakeConnection.
Username	Enter the user name to connect to the Snowflake account.
Password	Enter the password to connect to the Snowflake account.
Account	Enter the name of the Snowflake account.
Warehouse	Enter the Snowflake warehouse name.
Role	Specify the Snowflake role assigned to the user.
Additional JDBC URL Parameters	<p>Enter one or more JDBC connection parameters in the following format:</p> <p><param1>=<value>&<param2>=<value>&<param3>=<value>...</p> <p>For example:</p> <p>user=jon&warehouse=mywh&db=mydb&schema=public</p>

Creating a Snowflake Connection

Create a Snowflake connection before you create a Snowflake data object.

1. In the Developer tool, click **Window > Preferences**.
2. Select **Informatica > Connections**.
3. Expand the domain in the **Available Connections**.
4. Select the connection type **Database > SnowflakeConnection**, and click **Add**.
5. Enter a connection name and an optional description.
6. Select SnowflakeConnection as the connection type.
7. Click **Next**.
8. Configure the connection properties.
9. Click **Test Connection** to verify the connection to Snowflake.
10. Click **Finish**.

CHAPTER 4

PowerExchange for Snowflake Data Objects

This chapter includes the following topics:

- [Snowflake Data Object Overview, 14](#)
- [Snowflake Data Object Properties, 14](#)
- [Snowflake Data Object Read Operation, 15](#)
- [Snowflake Data Object Write Operation, 16](#)
- [Creating a Snowflake Data Object, 18](#)
- [Creating a Snowflake Data Object Operation, 18](#)
- [Rules and Guidelines for Snowflake Mappings, 18](#)

Snowflake Data Object Overview

A Snowflake data object is a physical data object that uses Snowflake as a source or target. A Snowflake data object is a physical data object that represents data based on a Snowflake resource.

You can configure the data object read and write operation properties that determine how you can read data from Snowflake sources and load data to Snowflake targets.

Create a Snowflake data object from the Developer tool. PowerExchange for Snowflake creates the data object read operation and data object write operation for the Snowflake data object. You can edit the advanced properties of the data object read or write operation and run a mapping.

Snowflake Data Object Properties

Specify the data object properties when you create the data object.

The following table describes the properties that you configure for the Snowflake data objects:

Property	Description
Name	Name of the Snowflake data object.
Location	The project or folder in the Model Repository Service where you want to store the Snowflake data object.
Connection	Name of the Snowflake connection.

Snowflake Data Object Read Operation

Create a mapping with a Snowflake data object read operation to read data from Snowflake.

Snowflake Data Object Read Operation Properties

Snowflake data object read operation properties include run-time properties that apply to the Snowflake data object.

The Developer tool displays advanced properties for the Snowflake data object operation in the Advanced view.

The following table describes the advanced properties for a Snowflake data object read operation:

Property	Description
Database	Overrides the database name specified in the connection.
Schema	Overrides the schema name specified in the connection.
Warehouse	Overrides the Snowflake warehouse name specified in the connection.
Role	Overrides the Snowflake user role specified in the connection. Note: The Data Integration Service ignores the role you specify in a source transformation.
Pre SQL	SQL statement that the Data Integration Service executes before extracting data from the source. For example, if you want to update records in the database before you extract the records from the table, specify a Pre-SQL statement.
Post SQL	SQL statement that the Data Integration Service executes after extracting data from the source. For example, if you want to delete some records after the latest records load, specify a Post-SQL statement.

Query Properties

Use the **Query** tab to specify the join, filter, and sort conditions.

You can configure the following query properties for Snowflake sources:

Select Distinct

Selects unique values from the source. The Data Integration Service filters out unnecessary data when you use the Snowflake data object in a mapping.

Join

The Join expression you specify in a read operation to join data from multiple sources in a Snowflake data object.

Select the sources on the **Data Object** tab and define the join condition.

You can use the native or platform expression to join specific records:

- When you configure a platform expression, in the **Tables** section, you can select the outer, inner, left, and right joins to join the tables.
- In the Relationship section, define the relationship for the join:
 - Left Field. The column on which you want to apply the join condition.
 - Operator. The operators that you can use to join tables. You can select one of the following operators: =, >, >=, <, <=, !=
 - Right Field. The value you specify to join the tables.
- When you configure a native expression, select the fields and define a join query syntax in the Join text field. You must specify only the condition and not the type of join in the query. The condition you specify in the text box for the native expression is appended to the join condition.

Filter

The filter value in a read operation. The filter specifies the WHERE clause of the SELECT statement. Use a filter to reduce the number of rows that the Data Integration Service reads from the Snowflake source. When you enter a source filter, the Integration Service adds a WHERE clause to the default query.

You can use the native or platform expression to filter specific records.

- Configure the platform expression, you can specify the filter using the following fields:
 - Left Field. The column on which you want to apply the filter condition.
 - Operator. The operators that you can use to filter tables. You can select one of the following operators: =, >, >=, <, <=, !=
 - Right Field. The value based on which you want to filter the records.
- When you configure a native expression, select the fields and define a filter query syntax in the Filter text field. The condition you specify in the text box for the native expression is appended to the filter condition.

Sort

Sorts the rows queried from the Snowflake source. The Data Integration Service adds the ports to the ORDER BY clause in the default query.

Snowflake Data Object Write Operation

Create a mapping to write data to Snowflake. Use the Snowflake connection, and define the write operation properties to write data to Snowflake.

You can perform insert, update, delete, and upsert operations on a Snowflake target.

Snowflake Data Object Write Operation Properties

Snowflake data object write operation properties include run-time properties that apply to the Snowflake data object.

The Developer tool displays advanced properties for the Snowflake data object operation in the Advanced view.

The following table describes the Advanced properties for a Snowflake data object write operation:

Property	Description
UpdateMode	Loads data to the target based on the mode you specify. Select one of the following modes: <ul style="list-style-type: none">- Update As Update. Updates all rows flagged for update.- Update Else Insert. Updates all rows flagged for update if they exist in the target and then inserts any remaining rows marked for insert.
Database	Overrides the database name specified in the connection.
Schema	Overrides the schema name specified in the connection.
Warehouse	Overrides the Snowflake warehouse name specified in the connection.
Role	Overrides the Snowflake user role specified in the connection. Note: The Data Integration Service ignores the role you specify in a target transformation.
Pre SQL	SQL statement that the Data Integration Service executes before extracting data from the source. For example, if you want to assign sequence object to a primary key field of the target table before you load data to the table, specify a Pre-SQL.
Post SQL	SQL statement that the Data Integration Service executes after extracting data from the source. For example, if you want to alter the table created by using create target option and assign constraints to the table before you load data to the table, specify a Post-SQL.
Batch Row Size	Number of rows that the Data Integration Service writes in a batch to the Snowflake target.
Number of local staging files	Enter the number of local staging files. The Data Integration Service writes data to the target after the specified number of local staging files are created.
Truncate Target Table	Truncates the database target table before inserting new rows. Select one of the following options: <ul style="list-style-type: none">- True. Truncates the target table before inserting all rows.- False. Inserts new rows without truncating the target table. Default is false.
Additional Write Runtime Parameters	Specify additional run-time parameters. For example: <code>remoteStage=CQA.CQA_SCHEMA.CQA_STAGE</code> Separate multiple runtime parameters with &.

Creating a Snowflake Data Object

Create a Snowflake data object to add to a mapping.

1. Select a project or folder in the **Object Explorer** view.
2. Click **File > New > Data Object**.
3. Select **Snowflake Data Object** and click **Next**.
The **Snowflake Data Object** dialog box appears.
4. Enter a name for the data object.
5. Click **Browse** next to the **Location** option and select the target project or folder.
6. Click **Browse** next to the **Connection** option and select the Snowflake connection from which you want to import the Snowflake object.
7. To add a resource, click **Add** next to the **Selected Resources** option.
The **Add Resource** dialog box appears.
8. Select the checkbox next to the Snowflake object you want to add and click **OK**.
9. Click **Finish**.
The data object appears under Data Objects in the project or folder in the **Object Explorer** view.

Creating a Snowflake Data Object Operation

You can create the data object read, write, or lookup operation for Snowflake data objects. You can add the Snowflake data object operation to a mapping.

1. Select the data object in the **Object Explorer** view.
2. Right-click and select **New > Data Object Operation**.
The **Data Object Operation** dialog box appears.
3. Enter a name for the data object operation.
4. Select the type of data object operation. You can choose to create a read or write operation.
5. Click **Add**.
The **Select Resources** dialog box appears.
6. Select the Snowflake data object for which you want to create the data object operation and click **OK**.
7. Click **Finish**.

The Developer tool creates the data object operation for the selected data object.

Rules and Guidelines for Snowflake Mappings

Use the following rules and guidelines when you create a mapping:

- Ensure that the source and target table names contain only uppercase letters.
- You cannot use the OR operator in a filter condition.

- You must define a primary key in the target table. If you do not define a primary key in the target table, the mapping fails to delete the record from or update the record in the target table.

CHAPTER 5

PowerExchange for Snowflake Mappings

This chapter includes the following topics:

- [PowerExchange for Snowflake Mappings Overview, 20](#)
- [Mapping Validation and Run-time Environments, 20](#)
- [Snowflake Mapping Example, 21](#)

PowerExchange for Snowflake Mappings Overview

After you create a Snowflake data object read or write operation, you can create a mapping to extract data from a Snowflake source or load data to a Snowflake target.

You can define properties in an operation to determine how the Data Integration Service must extract data from a Snowflake source or load data to a Snowflake target. You can extract data from one or more Snowflake sources, and load data to one or more Snowflake targets. When the Data Integration Service extracts data from the source or loads data to the target, it converts the data based on the data types associated with the source or the target.

Mapping Validation and Run-time Environments

You can validate and run mappings in the native environment or Spark engine.

The Data Integration Service validates whether the mapping can run in the selected environment. You must validate the mapping for an environment before you run the mapping in that environment.

Native environment

You can configure the mappings to run in the native or Hadoop environment. When you run mappings in the native environment, the Data Integration Service processes the mapping and runs the mapping from the Developer tool.

Spark Engine

When you run mappings on the Spark engine, the Data Integration Service pushes the mapping to a Hadoop cluster and processes the mapping on a Spark engine. The Data Integration Service generates an execution plan to run mappings on the Spark engine.

You can view the plan in the Developer tool before you run the mapping and in the Administrator tool after you run the mapping.

For more information about the Hadoop environment and Spark engines, see the *Informatica Big Data Management™ Administrator Guide*.

Snowflake Mapping Example

Your organization has a large amount of customer data from across regions stored in flat files. You organization needs to analyze data in the APAC region. Create a mapping that reads all the customer records from the flat file and write those records to a Snowflake table.

You can use the following objects in a Snowflake mapping:

Flat file input

The input file is a flat file that contains customer names and their details.

Create a flat file data object. Configure the flat file connection and specify the flat file that contains the customer data as a resource for the data object. Use the data object in a mapping as a read data object.

Transformations


Add a Filter transformation to extract customer data in the APAC region.

The Filter transformation filters the source data based on the value you specify for the region ID column. The Data Integration Service returns the rows that meet the filter condition.

Snowflake output

Create a Snowflake data object write operation. Configure the Snowflake connection and specify the Snowflake object as a target for the data object. Use the data object in a mapping as a target data object.

The following image shows the Snowflake mapping example:



When you run the mapping, the Data Integration Server reads customer records from the flat file and writes to the Snowflake table.

APPENDIX A

Snowflake Data Type Reference

This appendix includes the following topics:

- [Data Type Reference Overview, 22](#)
- [Snowflake and Transformation Data Types, 22](#)

Data Type Reference Overview

Developer Tool uses the following data types in Snowflake mappings:

- Snowflake native data types. Snowflake data types appear in Snowflake definitions in a mapping.
- 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 Data Integration Service uses to move data across platforms. They appear in all transformations in a mapping.

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

Snowflake and Transformation Data Types

The following table lists the Snowflake data types that Developer Tool supports and the corresponding transformation data types:

Snowflake Data Type	Transformation Data Type	Range and Description
BINARY (VARBINARY)	Binary	Maximum value: 8,388,60 Default value is 8,388,60.
BOOLEAN	String	A Boolean attribute.
DATE	Datetime	Date and time values.

Snowflake Data Type	Transformation Data Type	Range and Description
FLOAT (DOUBLE, DOUBLE PRECISION, REAL, FLOAT, FLOAT4, FLOAT8)	Double	Floating point numbers with double-precision (64 bit). Maximum value: 1.7976931348623158e+307 Minimum value: -1.79769313486231E+307
NUMBER (DECIMAL, NUMERIC, INT, INTEGER, BIGINT, SMALLINT, TINYINT, BYTEINT)	Decimal	Number with 28-bit precision and scale. Note: PowerExchange for Snowflake does not support Decimal values of precision that exceeds 28-bit for the source or target because of an SDK limitation.
TIMESTAMP_LTZ	Datetime	Date and time values.
TIMESTAMP_NTZ (TIMESTAMP_NTZ, datetime)	Datetime	Date and time values.
TIMESTAMP_TZ	Datetime	Date and time values.
VARCHAR (TEXT, CHAR, CHARACTER, STRING)	String	Maximum value: 16,777,216 Default value is 16,777,216.

INDEX

C

- client component
 - installation [10](#)
- create
 - data object operation
 - create [18](#)
 - Snowflake data object [18](#)
- creating
 - Snowflake connection [13](#)

D

- data object read operation
 - properties [15](#)
- data object write operation
 - properties [17](#)
- data types [22](#)

I

- installation
 - client component [10](#)
 - installation on UNIX
 - server component [9](#)

M

- mapping
 - example [21](#)

N

- native environment
 - mappings [20](#)

P

- PowerExchange for Snowflake
 - configuration [9](#)

- PowerExchange for Snowflake (*continued*)

- EBF [7](#)
 - example [7](#)
 - installation [9](#)
 - overview [7](#)
- PowerExchange for Snowflake mappings
 - overview [20](#)
- properties
 - data object read operation [15](#)
 - data object write operation [17](#)

S

- server component
 - installation on UNIX [9](#)
- snowflake
 - architecture [7](#)
 - cloud data warehouse [7](#)
- Snowflake
 - data object properties [14](#)
 - data object read operation [15](#)
 - data object write operation [16](#)
- Snowflake connection
 - properties [12](#)
- Snowflake connections
 - creating [13](#)
 - overview [12](#)
- Snowflake data object
 - create [18](#)
 - overview [14](#)
- Snowflake data types [22](#)
- Snowflake read operation properties
 - configuring filter conditions [15](#)
 - configuring join conditions [15](#)
 - configuring sort conditions [15](#)
 - query [15](#)
- Spark engine
 - mappings [20](#)

T

- transformation data types [22](#)