



Informatica® PowerExchange for Netezza
10.5

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, 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 Matrices.	6
Informatica Velocity.	6
Informatica Marketplace.	6
Informatica Global Customer Support.	6
 Chapter 1: Introduction to PowerExchange for Netezza.....	 7
PowerExchange for Netezza Overview.	7
Introduction to Netezza.	7
 Chapter 2: PowerExchange for Netezza Configuration.....	 8
PowerExchange for Netezza Configuration Overview.	8
Prerequisites.	8
Configure the Netezza ODBC Driver	9
 Chapter 3: Netezza Connections.....	 11
Netezza Connection Overview.	11
Netezza Connection Properties.	11
Creating a Netezza Connection.	12
SSL Configuration.	13
Configuring SSL Authentication on Linux.	13
Configuring SSL Authentication on Windows.	13
 Chapter 4: Netezza Data Objects and Mappings.....	 14
Netezza Data Objects and Mappings Overview.	14
Netezza Data Object Properties.	15
Netezza Data Object Read Operation Properties.	15
Output Properties of a Netezza Data Object Read Operation.	15
Ports Properties.	16
Sources Properties.	16
Query Properties.	16
Run-time Properties.	17
Advanced Properties.	17
Netezza Data Object Write Operation Properties.	19
Target Properties of a Netezza Data Object Write Operation.	19
Input Properties of a Netezza Data Object Write Operation.	19

Ports Properties.	19
Run-time Properties.	19
Advanced Properties.	20
Importing a Netezza Data Object.	23
Creating a Netezza Data Object Operation.	24
Netezza Mappings.	24
Chapter 5: Netezza Run-time Processing Overview.	25
Netezza Run-time Processing Overview.	25
Partitioning.	25
Parameterization for Netezza Sources.	26
Parameterization for Netezza Targets.	26
Rules and Guidelines for Dynamic Mappings.	27
Appendix A: Data Type Reference.	28
Data Type Reference Overview.	28
Netezza and Transformation Data Types.	28
Index.	30

Preface

Use the *Informatica® PowerExchange® for Netezza User Guide* to learn how to read from or write to Netezza by using the Developer tool. Learn to create a Netezza connection and develop and run mappings in an Informatica domain. It also includes information on partitioning, dynamic mapping, and parameterization of Netezza sources and targets.

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 Network

The Informatica Network is the gateway to many resources, including the Informatica Knowledge Base and Informatica Global Customer Support. To enter the Informatica Network, visit <https://network.informatica.com>.

As an Informatica Network member, you have the following options:

- Search the Knowledge Base for product resources.
- View product availability information.
- Create and 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 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 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 Product Availability Matrices

Product Availability Matrices (PAMs) indicate the versions of the operating systems, databases, and types of data sources and targets that a product release supports. You can browse the Informatica 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 and based on real-world experiences from hundreds of data management projects. Informatica Velocity represents the collective knowledge of Informatica consultants who work with organizations around the world to plan, develop, deploy, and maintain successful data management solutions.

You can find 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 extend and enhance your Informatica implementations. Leverage any of the hundreds of solutions from Informatica developers and partners on the Marketplace to improve your productivity and speed up time to implementation on your projects. You can find the Informatica Marketplace at <https://marketplace.informatica.com>.

Informatica Global Customer Support

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

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

<https://www.informatica.com/services-and-training/customer-success-services/contact-us.html>.

To find online support resources on the Informatica Network, visit <https://network.informatica.com> and select the eSupport option.

CHAPTER 1

Introduction to PowerExchange for Netezza

This chapter includes the following topics:

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

PowerExchange for Netezza Overview

You can use PowerExchange for Netezza to connect to a Netezza database and read and write data in bulk. You can use PowerExchange for Netezza to process large volumes of data.

The Data Integration Service reads and writes Netezza data through a Netezza external table. An external table definition is stored within the Netezza database but the data is saved externally in a location that is accessible to the Netezza host or the client system.

To read data from a Netezza database, the Data Integration Service creates an external table and named pipe in the pipe directory path that you specify for data extraction. It then reads the data from the named pipe.

To write data to a Netezza database, the Data Integration Service creates an external table in the pipe directory path that you specify for data loading. It writes the data first to the external table and then to a staging table. The Data Integration Service then writes the data from the staging table to the Netezza target.

You can run Netezza mappings in the native environment.

Example

Your organization is a retail chain with sales outlets across multiple countries. You store the historical sales information in a Netezza database. You can use PowerExchange for Netezza to read sales information from the Netezza database and use it for business analysis and decision making.

Introduction to Netezza

Netezza is a data warehousing appliance that organizations use to store large volumes of data. Netezza provides high-performance analytics and fast querying capabilities.

CHAPTER 2

PowerExchange for Netezza Configuration

This chapter includes the following topics:

- [PowerExchange for Netezza Configuration Overview, 8](#)
- [Prerequisites, 8](#)
- [Configure the Netezza ODBC Driver , 9](#)

PowerExchange for Netezza Configuration Overview

PowerExchange for Netezza installs with the Informatica services and clients.

To configure PowerExchange for Netezza, complete the prerequisites.

Prerequisites

Before you use PowerExchange for Netezza, you must complete the following prerequisites:

- Install and configure the Informatica services.
- Install and configure the Developer tool. You can install the Developer tool when you install Informatica clients.
- Create a Data Integration Service and a Model Repository Service in the Informatica domain.
- Verify that the Netezza database user has the following privileges on the database:
 - CREATE EXTERNAL TABLE
 - CREATE TABLE
 - DELETE
 - DROP
 - INSERT
 - LIST
 - SELECT

- TRUNCATE
- UPDATE
- Download the Netezza JDBC driver from the IBM website and install the driver.
The Developer tool uses the Netezza JDBC driver to connect to the Netezza database.
- To import metadata from Netezza tables, copy the Netezza JDBC jars to the following directory on the node where you installed the Developer tool:
`<Informatica installation directory>/clients/externaljdbcjars`
- Download the Netezza ODBC driver from the IBM website and install the driver.
The Data Integration Service uses the Netezza ODBC driver to run Netezza mappings.
- Use the Netezza ODBC driver to create an ODBC data source for the Netezza database that you want to connect to.
The ODBC data source name is the connection string that you must enter while creating a Netezza connection.
- Install the Netezza client drivers on the machine where you installed the Informatica services.

Configure the Netezza ODBC Driver

On Windows, you must install the Netezza ODBC driver. Create a DSN using the ODBC Data Source Administrator.

On Linux, you must set the required environmental variables for the ODBC driver and create the DSN entries.

1. Add the path of the `odbc.ini` file to the ODBCINI environment variable.
For example: `setenv ODBCINI "/data/home/qamercury/ODBCINI/odbc.ini"`
2. To set the ODBCINST environment variable to the location of the `odbcinst.ini` file, use the following syntax:
For example: `setenv ODBCINST /data/home/qamercury/ODBCINI/odbcinst.ini`
3. To set the NZ_ODBC_INI_PATH environment variable to the location of the `odbc.ini` file, use the following syntax:
For example: `setenv NZ_ODBC_INI_PATH /data/home/qamercury/ODBCINI`
4. To set the LD_LIBRARY_PATH environment variable to the location that contains the Netezza libraries, use the following syntax:
For example: `setenv LD_LIBRARY_PATH ".:./export/qa_adp/thirdparty/netezza/linux.64/lib64:$LD_LIBRARY_PATH"`
5. To set the \$ODBCINI_HOME environment variable, use the following syntax:
For example, `setenv ODBCINI_HOME /data/home/qamercury/ODBCINI`
6. Edit the `odbc.ini` file.
You can find the `odbc.ini` file in the following directory:
`$ cp $ODBCHOME/odbc.ini $ODBCINI_HOME/odbc.ini`
Add an entry for the Netezza data source under the section [ODBC Data Sources] and configure the data source.

For example:

```
[Sample Netezza ODBC DSN]
Driver=/export/qa_adp/thirdparty/netezza/linux.64/lib64/libnzodbc.so
Description=NetezzaSQL ODBC
Servername=adaptersnz2.informatica.com
Port=5480
Database=ADPQA_DB
Username=adpqa_
Password=adpqa
StripCRLF=false
ReadOnly=false
ShowSystemTables=false
DateFormat=1
NumericAsChar=false
DebugLogging=true
```

For more information about Netezza connectivity, see the Netezza ODBC driver documentation.

7. Restart the Informatica services.

CHAPTER 3

Netezza Connections

This chapter includes the following topics:

- [Netezza Connection Overview, 11](#)
- [Netezza Connection Properties, 11](#)
- [Creating a Netezza Connection, 12](#)
- [SSL Configuration, 13](#)

Netezza Connection Overview

Use a Netezza connection to access a Netezza database.

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

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

Netezza Connection Properties

Use a Netezza connection to access a Netezza database. The Netezza connection is a database connection. You can create and manage a Netezza connection in the Administrator tool or the Developer tool.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Netezza connection properties:

Property	Description
Name	Name of the connection. The name is not case sensitive and must be unique within the domain. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
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.

Property	Description
Description	Description of the connection. The description cannot exceed 765 characters.
Location	Domain where you want to create the connection.
Type	Connection type. Select Netezza .
User name	User name with the appropriate permissions to access the Netezza database.
Password	Password for the database user name.
JDBC Url	JDBC URL that the Developer tool must use when it connects to the Netezza database. Use the following format: <code>jdbc:netezza://<hostname>:<port>/<database name></code>
Connection String	Name of the ODBC data source that you want to use to connect to the Netezza database.
Timeout	Number of seconds that the Developer tool waits for a response from the Netezza database before it closes the connection.

Creating a Netezza Connection

Before you import Netezza data objects or run mappings, create a Netezza connection.

1. Click **Window > Preferences**.
2. Select **Informatica > Connections**.
3. Expand the domain in the **Available Connections** list.
4. Select **Database > Netezza** and then click **Add**.
5. Enter a connection name.
6. Optionally, enter a connection ID and description.
7. Select the domain where you want to create the connection.
8. Select the connection type as **Netezza**.
9. Click **Next**.
10. Configure the connection properties.
11. Click **Test Connection** to verify the connection to the Netezza database.
12. Click **OK**.

SSL Configuration

You can use the Secure Sockets Layer (SSL) protocol to configure a secure connection between Netezza clients and the Netezza server.

The Data Integration Service uses a trust certificate file to establish a secure connection with the Netezza server over SSL. To establish a secure connection with the Netezza server over SSL, define the security type in the ODBC settings and specify the path to the cacert file.

Configuring SSL Authentication on Linux

You can configure SSL authentication for the Netezza server through ODBC by using the cacert trust certificate file.

1. Open the `odbc.ini` file and add entries for the security level and cacert file in the NetezzaSQL ODBC section.
2. Add the `securityLevel=onlySecured` attribute. The `securityLevel` determines the method that the driver uses to encrypt the data sent between the driver and the database server.
3. Add the `caCertFile=<file path>` attribute. The `caCertFile` is the location and name of the trust certificate file.

Configuring SSL Authentication on Windows

1. Open the **Administrative Tools** from the Windows Control Panel.
2. Open the **Data Sources (ODBC)** shortcut.
The ODBC Data Source Administrator appears.
3. Click the **System DSN** tab.
4. Select **IBM Netezza ODBC Driver** and click **Configure**.
5. Click the **SSL DSN Options** tab.
6. Under **Security Level**, select **Only Secured**.
7. In the **CA Certificate File** field, specify the path to the cacert trust certificate and click **OK**.

CHAPTER 4

Netezza Data Objects and Mappings

This chapter includes the following topics:

- [Netezza Data Objects and Mappings Overview, 14](#)
- [Netezza Data Object Properties, 15](#)
- [Netezza Data Object Read Operation Properties, 15](#)
- [Output Properties of a Netezza Data Object Read Operation, 15](#)
- [Netezza Data Object Write Operation Properties, 19](#)
- [Target Properties of a Netezza Data Object Write Operation, 19](#)
- [Input Properties of a Netezza Data Object Write Operation, 19](#)
- [Importing a Netezza Data Object, 23](#)
- [Creating a Netezza Data Object Operation, 24](#)
- [Netezza Mappings, 24](#)

Netezza Data Objects and Mappings Overview

A Netezza data object is a physical data object that uses a Netezza table as a source or target. A Netezza data object is the representation of data that is based on a Netezza table.

To create a Netezza data object, import metadata from a Netezza table into the Developer tool. Create a data object read operation or a data object write operation based on the Netezza data object.

To read data from the Netezza table, create a data object read operation based on the Netezza data object. Configure the read operation properties to determine how the Data Integration Service must read data from the Netezza table. Add the read operation as a Read transformation in a mapping.

To write data to the Netezza table, create a data object write operation based on the Netezza data object. Configure the write operation properties to determine how the Data Integration Service must write data to the Netezza table. Add the write operation as a Write transformation in a mapping.

Netezza Data Object Properties

The Netezza **Overview** view displays general information about the Netezza data object and the object properties that apply to the Netezza table that you import.

You can configure the following properties for a Netezza data object:

General Properties

You can configure the following general properties for a Netezza data object:

- Name. Name of the Netezza data object.
- Description. Description of the Netezza data object.
- Connection. Name of the Netezza connection.

Object Properties

You can configure the following general properties and column properties for the Netezza table that you add in the data object:

- Name. Business name of the Netezza table.
- Description. Description of the Netezza table.
- Native name. Native name of the Netezza table including the entity in which the table exists.
- Path Information. Path to the Netezza table.
- Column Properties. Name, native name, data type, precision, scale, and description of the columns in the Netezza table. You can also define the primary key and specify if a column can contain null values.

Netezza Data Object Read Operation Properties

The Data Integration Service reads data from a Netezza table based on the data object read operation properties that you specify.

When you create a data object read operation, the Developer tool creates a Source transformation and an Output transformation.

The Source transformation represents the data that the Data Integration Service reads from the Netezza table. You can configure the Source transformation to select distinct rows from the Netezza table and sort data.

The Output transformation represents the data that the Data Integration Service passes into the mapping pipeline. Select the Output transformation to edit the ports, sources, query, run-time, and advanced properties.

Output Properties of a Netezza Data Object Read Operation

The Output transformation defines the run-time properties that the Data Integration Service uses to read data from the Netezza table.

You can configure the Output transformation to perform the following tasks:

- Edit the ports properties.
- Add or remove sources.
- Define a filter or sort condition.
- Change or parameterize the Netezza connection.
- Define the advanced properties that the Data Integration Service must use to read data.

Ports Properties

The **Ports** tab lists the name, data type, precision, scale, and description for all the ports that the data object read operation contains.

You can configure the following ports properties in the data object read operation:

Property	Description
Name	Name of the port.
Type	Data type of the port.
Precision	Maximum number of digits for numeric data types or maximum number of characters for string data types. The precision includes the scale for numeric data types.
Scale	Maximum number of digits after the decimal point for numeric values.
Description	Description of the port.

Sources Properties

Use the **Sources** tab to update the list of Netezza tables from which you want to read data. You can read data from multiple Netezza tables by using the same data object.

Query Properties

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

The following table describes the query properties that you can configure for Netezza sources:

Property	Description
Select Distinct	Selects unique values from the source. The Data Integration Service filters out unnecessary data when you use the Netezza data object in a mapping.
Join	User-defined join in a Netezza data object. A user-defined join specifies the condition used to join data from multiple sources in the same Netezza data object.
Filter	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 Netezza source. When you enter a source filter, the Developer tool adds a WHERE clause to the default query.
Sort	Sorts the rows queried from the Netezza source. The Data Integration Service adds the ports to the ORDER BY clause in the default query.

Platform Expression

You can use the platform filter expression to select specific records from Netezza sources based on the filter condition that you specify.

The following table describes the properties that you can specify when you use the platform expression filter:

Property	Description
Expression Type	Type of filter expression that you want to use to filter records. Default is Platform Expression.
Left Field	Column on which you want to apply the filter condition.
Operator	Simple operators that you can use to filter records. You can select one of the following operators: =, !=, <, <=, >, >=
Right Field	Value based on which you want to filter the records.

Run-time Properties

The **Run-time** tab displays the name of the connection that the Data Integration Service uses to read data from the Netezza table. You can select a different connection or you can parameterize the connection. You can also configure partitioning.

Advanced Properties

The **Advanced** tab displays the advanced properties that the Data Integration Service uses to read data from the Netezza table.

You can configure the following advanced properties in the data object read operation:

Socket Buffer Size

Buffer size that the Data Integration Service uses to read data.

Enter a value between 4096 and 2147483648 bytes.

Default is 8388608 bytes.

You can set the socket buffer size to 25% to 50% of the DTM buffer size to increase performance. You might need to test different settings for optimal performance.

Pipe Directory Path

Path where the Data Integration Service creates a named pipe for the external table. The Data Integration Service reads data from the named pipe.

Enter a path that does not use an NFS mount.

If you do not specify the path, the Data Integration Service creates the pipe in the following directory:

```
<Informatica installation directory>/server/bin
```

Delimiter

Character that separates successive input fields.

Set the delimiter to any value that the Netezza Performance Server supports. The value can be a part of the data that the Netezza source contains.

Default is |.

Escape Character

Escape character of the external table.

If the data contains NULL, CR, and LF characters in the Char or Varchar fields, you must add the escape character in the source data before reading data. Enter the escape character before the data.

Enter a backslash (\) as the escape character.

Null Value

NullValue parameter of the external table.

The string that represents a null value for the column. If you set this parameter as 'a', the Data Integration Service treats 'a' data as NULL.

Maximum value is one character.

Default is ".

Synchronize Netezza Null

Select this option if you want the Data Integration Service to treat null values from the Netezza source as null values in Informatica Developer.

Default is not selected.

SQL Query

Overrides the default query.

Enter the SQL query that the Data Integration Service must use to query data from the Netezza source.

Pre SQL

SQL query that the Data Integration Service must run before it reads data from the Netezza source.

Post SQL

SQL query that the Data Integration Service must run after it reads data from the Netezza source.

Netezza Custom Properties

Custom properties that you want to pass to the external table.

Enter any parameter that the external table supports along with the parameter value

Use the following format:

```
<parameter name> <value>
```

To enter multiple parameter name and value pairs, separate them with a space character.

Source Table Name

Overrides the source table name.

For example, you can parameterize the source table name to override the table name in the mapping.

Source Table Prefix

Name of the schema from which the Data Integration Service must read data.

Netezza Data Object Write Operation Properties

The Data Integration Service writes data to a Netezza table based on the data object write operation properties that you specify.

Target Properties of a Netezza Data Object Write Operation

The target properties represent the data that the Data Integration Service writes to the Netezza table. Select the target properties to view data such as the name and description of the Netezza table, and the column properties.

Input Properties of a Netezza Data Object Write Operation

The Input transformation defines the run-time properties that the Data Integration Service uses to write data to the Netezza table.

In the Input transformation, you can also edit the port properties, and change or parameterize the Netezza connection.

Ports Properties

The **Ports** tab lists the name, data type, precision, scale, and description for all the ports that the data object write operation contains.

You can configure the following ports properties in the data object write operation:

Property	Description
Name	Name of the port.
Type	Data type of the port.
Precision	Maximum number of digits for numeric data types or maximum number of characters for string data types. The precision includes the scale for numeric data types.
Scale	Maximum number of digits after the decimal point for numeric values.
Description	Description of the port.

Run-time Properties

The **Run-time** tab displays the name of the connection that the Data Integration Service uses to write data to the Netezza table. You can select a different connection or you can parameterize the connection. You can also configure partitioning.

Advanced Properties

The **Advanced** tab displays the advanced properties that the Data Integration Service uses to write data to the Netezza table.

You can configure the following advanced properties in the data object write operation:

Pre SQL

SQL query that the Data Integration Service must run before it writes data to the Netezza target.

For example, you can run an SQL query to drop indexes before the Data Integration Service writes data to the Netezza target.

Post SQL

SQL query that the Data Integration Service must run after it writes data to the Netezza target.

For example, you can run an SQL query to re-create indexes after the Data Integration Service writes data to the Netezza target.

Table Name Prefix

Prefix that the Data Integration Service must add to the target table name.

For example, you can enter a target owner name to override the table name prefix in the mapping.

Target Table Name

Overrides the target table name.

For example, you can parameterize the target table name to override the table name in the mapping.

Insert

The Data Integration Service inserts rows into the target table and enforces key constraints.

Default is selected.

Delete

The Data Integration Service deletes rows from the target table.

Default is selected.

Update

The Data Integration Service updates the rows based on the update strategy option that you specify.

Select one of the following values:

- Update as Update. The Data Integration Service updates all rows flagged for update.
- Update as Insert. The Data Integration Service inserts all rows flagged for update.
- Update else Insert. The Data Integration Service first updates all rows flagged for update if they exist in the target. It then inserts the remaining rows marked for insert.
- None. The Data Integration Service does not update any row.

Default is None.

When you use the Update property, the Data Integration Service does not enforce key constraints and writes duplicate rows with the same primary key into the target table.

Truncate Target Table Option

The Data Integration Service truncates the target before it writes data to the target.

Default is selected.

If you specify an SQL statement in the **Pre SQL** property, the Data Integration Service runs the SQL statement before it truncates the table.

Delimiter

Character that separates successive input fields.

Set the delimiter to any value that the Netezza Performance Server supports. The value must not be a part of the input data.

Default is |.

Null Value

NullValue parameter of the external table. The string that represents a null value for the column. If you set this parameter as 'a', the Data Integration Service treats 'a' data as NULL.

Maximum value is one character.

Default is "".

Synchronize Platform Null

Select this option if you want the Data Integration Service to write null values from Informatica Developer as null values into the Netezza target.

Default is not selected.

Escape Character

Escape character of the external table. If the data contains NULL, CR, and LF characters in the Char or Varchar fields, you must add an escape character for these fields before writing data to the target.

Enter a backslash (\) as the escape character.

Quoted Value

QUOTEDVALUE parameter of the external table. The quoted value is not a part of the data.

Select one of the following values:

- NO. Select this option to omit the quotes.
- SINGLE. Select this option to enclose the column value in single quotes.
- DOUBLE. Select this option to enclose the column value in double quotes.

Default is NO.

Enable Key Constraint

This option is applicable only when you want to insert rows. When you select this option, the Data Integration Service enforces key constraints and does not write duplicate rows into the target table.

If you clear this option, the Data Integration Service does not enforce key constraints and writes duplicate rows with the same primary key to the target.

Default is not selected.

Enable Advanced Statistics

Select this option to write advanced statistics in the workflow log such as the number of duplicate rows and the actual number of rows that the Data Integration Service wrote to the target.

For example, a Netezza target contains 1000 rows and there are 100 duplicate rows. You configure the Data Integration Service to enforce key constraints. When you run a mapping to write data to the target, the Data Integration Service writes only 900 rows to the target and rejects 100 rows. When you select the

Enable Advanced Statistics option, the workflow log indicates that there were 100 duplicate rows and the Data Integration Service wrote only 900 rows to the target.

Default is not selected.

Duplicate Row Handling

Determines how the Data Integration Service handles duplicate rows.

Select one of the following values:

- **FIRST.** The Data Integration Service passes the first row to the target and rejects the rows that follow with the same primary key.
- **LAST.** The Data Integration Service passes the last duplicate row to the target and discards the rest of the rows.

Default is FIRST.

Socket Buffer Size

Buffer size that the Data Integration Service uses to write data.

Enter a value between 4096 and 2147483648 bytes.

Default is 8388608 bytes.

Set the socket buffer size to 25% to 50% of the DTM buffer size to increase session performance. You might need to test different settings for optimal performance.

Control Character

CTRLCHARS parameter of the external table to transfer data that contains control characters. You can enter control characters for Char and Varchar fields. If you enter a control character, you must add an escape character for the NULL, CR, and LF fields.

Default is TRUE.

Crinstring

CRINSTRING parameter to transfer data that contains carriage returns (CR). You can enter a non escape CR in Char or Varchar fields. To write the control characters present in the Char and Varchar fields, set the CTRLCHARS and CRINSTRING parameters to TRUE in the Netezza table properties.

Default is TRUE.

Pipe Directory Path

Path where the Data Integration Service creates the pipe for the external table.

Enter a path that does not use an NFS mount.

If you do not specify the path, the Data Integration Service uses the following directory to create the pipe for the external table:

`<Informatica installation directory>/server/bin`

Error Log Directory Name

Directory where Netezza creates an error log if there are data errors. The error log contains the error messages.

By default, Netezza creates the error log in the following directory on the machine that hosts the Netezza Performance Server:

`/tmp`

Bad File Name

Directory where the Data Integration Service creates a bad file. The bad file contains the rejected records.

By default, the Data Integration Service creates a bad file in the following directory:

```
$PMBadFileDir
```

Error Threshold

Number of non-fatal errors that the Data Integration Service can encounter before it fails the mapping.

If you specify 0, non-fatal errors do not cause the mapping to fail.

Default is 1.

Custom Properties

Custom properties that you want to pass to the external table.

Enter any parameter that the external table supports along with the parameter value.

Use the following format:

```
<parameter name> <value>
```

To enter multiple parameter name and value pairs, separate them with a space character.

Target Schema Strategy

The target schema strategy for the target object. You can select from the following options:

- **RETAIN.** The Data Integration Service retains the existing target schema.
- **CREATE.** The Data Integration Service creates or replace the Netezza target at run time when you run a dynamic mapping.
- **Assign Parameter.** Specify the target schema strategy options as a parameter value.

Importing a Netezza Data Object

Import a Netezza data object to specify the Netezza table from which you want to read data or write data to. You can then create a data object operation based on the data object.

1. Select a project or folder in the **Object Explorer** view.
2. Click **File > New > Data Object**.
3. Select **Netezza Data Object** and click **Next**.
The **New Netezza 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 Netezza connection from which you want to import the Netezza table metadata.
7. To add a Netezza table, click **Add** next to the **Selected Resources** option.
The **Add Resource** dialog box appears.
8. Select the Netezza table that you want to add.

You can also search for a table by entering the table name in the **Name Filter** field.

9. Click **Finish**.

The data object appears under Physical Data Objects in the project or folder in the **Object Explorer** view.

Creating a Netezza Data Object Operation

Create a Netezza data object operation from a Netezza data object. You can create a read operation or a write operation. You can then add the read operation as a source in a mapping and add the write operation as a target in a mapping.

1. Select the Netezza data object in the **Object Explorer** view.
2. Right-click the data object and select **New > Data Object Operation**.
The **Data Object Operation** dialog box appears.
3. Enter a name for the data object operation.
4. Specify the type of data object operation that you want to create.
 - To create a data object read operation, select **Read** from the **Capabilities** list.
 - To create a data object write operation, select **Write** from the **Capabilities** list.
5. Click **Add**.
The **Select Resources** dialog box appears.
6. Select the Netezza table 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 Netezza data object.

Netezza Mappings

After you create a Netezza data object operation, you can create a mapping to read data from or write data to Netezza.

You can configure mappings to run in the native environment. When you run mappings in the native environment, the Data Integration Service processes the mapping. You can also configure dynamic mappings with Netezza sources and targets and run the mappings in the native environment.

To read data from a Netezza database, add the data object read operation as a Read transformation in the mapping. To write data to a Netezza database, add the data object write operation as a Write transformation in the mapping.

Validate and run the mapping. You can deploy the mapping and run it, or you can add the mapping to a Mapping task in a workflow and run the workflow.

CHAPTER 5

Netezza Run-time Processing Overview

This chapter includes the following topics:

- [Netezza Run-time Processing Overview, 25](#)
- [Partitioning, 25](#)
- [Parameterization for Netezza Sources, 26](#)
- [Parameterization for Netezza Targets, 26](#)
- [Rules and Guidelines for Dynamic Mappings, 27](#)

Netezza Run-time Processing Overview

When you create a Netezza data object read or write operation, you define properties that determine how the Data Integration Service reads data from or writes data to a Netezza database.

You can configure partitioning and parameterization in the run-time properties.

Partitioning

You can configure partitioning for Netezza mappings that you run in the native environment. When a mapping that is enabled for partitioning contains a Netezza data object operation as a Read transformation or Write transformation, the Data Integration Service can use multiple threads to read and write data.

Some transformations do not support partitioning. When a mapping enabled for partitioning contains a transformation that does not support partitioning, the Data Integration Service uses one thread to run the transformation. The Data Integration Service can use multiple threads to run the remaining mapping pipeline stages.

You can configure dynamic partitioning for Netezza data objects. When you configure dynamic partitioning, the Data Integration Service determines the number of partitions to create at run time.

To configure dynamic partitioning, open the Netezza data object read or write operation, and select the **Dynamic Partitioning** option in the **Run-time** tab.

After you configure dynamic partitioning, you can configure the maximum parallelism.

Parameterization for Netezza Sources

You can parameterize the Netezza connection and data object read operation properties to override the mapping properties at run time.

You can parameterize the following read operation properties for a Netezza source:

- Socket Buffer Size
- Pipe Directory Path
- Delimiter
- Escape Character
- Null Value
- SQL Query
- Pre SQL
- Post SQL
- Netezza Custom Properties
- Source Table Name
- Source Table Prefix

Parameterization for Netezza Targets

You can parameterize the Netezza connection and data object write operation properties to override the mapping properties at run time.

You can parameterize the following write operation properties for a Netezza target:

- Pre SQL
- Post SQL
- Table Name Prefix
- Target Table Name
- Update
- Delimiter
- Null Value
- Escape Character
- Quoted Value
- Duplicate Row Handling
- Socket Buffer Size
- Control Character
- Crinstring
- Pipe Directory Path
- Error Log Directory Name
- Bad File Name
- Netezza Custom Properties

Rules and Guidelines for Dynamic Mappings

Consider the following rules and guidelines when you create or replace the Teradata target using a dynamic mapping:

- When you select the **Create - Create or replace table at run time** option for the **Target Schema Strategy** property and run a dynamic mapping, the Integration Service drops the target table and recreates it. When you rerun a mapping and you want to retain the target table that resulted from the first mapping, choose the **RETAIN - Retain existing target schema** option before you run the mapping.
- You do not get the primary key information from the source.
- The Data Integration Service creates all target fields as nullable fields even if the source fields are not null.
- The Data Integration Service converts the native data types from the source to the following data types in the target:

Source Data Type	Target Data Type
byteint	integer
smallint	integer
float	double
char	nvarchar
nchar	nvarchar
date	timestamp
time	timestamp
boolean	numeric

APPENDIX A

Data Type Reference

This appendix includes the following topics:

- [Data Type Reference Overview, 28](#)
- [Netezza and Transformation Data Types, 28](#)

Data Type Reference Overview

Informatica Developer uses the following data types in Netezza mappings:

- Netezza native data types. Netezza data types appear in the column properties of the physical data object.
- 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. Transformation data types 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.

Netezza and Transformation Data Types

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

Netezza Data Type	Transformation Data Type	Range and Description
BigInt	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Bool	String	1 character
ByteInt	Small Integer	Precision 5, scale 0
Char	String	1 to 104,857,600 characters

Netezza Data Type	Transformation Data Type	Range and Description
Date	Date/Time	Jan. 1, 0001 A.D. to Dec. 31, 9999 A.D. (precision to the nanosecond)
Decimal	Decimal if the precision is 25 or lesser Double if the precision is greater than 25	For transformations that support precision up to 38 digits, the precision is 1 to 38 digits, and the scale is 0 to 38. For transformations that support precision up to 28 digits, the precision is 1 to 28 digits, and the scale is 0 to 28. If you specify the precision greater than the maximum number of digits, the Data Integration Service converts decimal values to double in high precision mode. Double: Precision 15
Float4	Double	Precision 15
Float8	Double	Precision 15
Integer	Integer	-2,147,483,648 to 2,147,483,647 Precision 10, scale 0
NChar(m)	nString	1 to 104,857,600 characters
Numeric	Decimal if the precision is 25 or lesser Double if the precision is greater than 25	For transformations that support precision up to 38 digits, the precision is 1 to 38 digits, and the scale is 0 to 38. For transformations that support precision up to 28 digits, the precision is 1 to 28 digits, and the scale is 0 to 28. If you specify the precision greater than the maximum number of digits, the Data Integration Service converts decimal values to double in high precision mode. Double: Precision 15
NVarchar(m)	nString	1 to 104,857,600 characters
Real	Real	Double-precision floating-point numeric value Precision 7, scale 0
SmallInt	Small Integer	Precision 5, scale 0
Time	Date/Time	Jan. 1, 0001 A.D. to Dec. 31, 9999 A.D. (precision to the microsecond)
Timestamp	Date/Time	January 1, 0001 00:00:00.000000 to December 31, 9999 23:59:59.999999 (precision to the microsecond)
Varchar	String	1 to 104,857,600 characters

INDEX

N

Netezza connections

- creating [12](#)
- overview [11](#)
- properties [11](#)

Netezza data objects

- creating [23](#)
- overview [14](#)
- properties [15](#)

Netezza data type

- overview [28](#)

Netezza data types

- mapping with transformation data types [28](#)

Netezza database

- overview [7](#)

Netezza mappings

- overview [24](#)

Netezza parameterization

- for sources [26](#)
- for targets [26](#)

Netezza read operation

- creating [24](#)
- output transformation [15](#)
- source transformation [15](#)

Netezza read operation properties

- advanced [17](#)

Netezza read operation properties (*continued*)

- configuring filter conditions [16](#)
- configuring join conditions [16](#)
- configuring sort conditions [16](#)
- ports [16](#)
- query [16](#)
- run-time [17](#)
- sources [16](#)

Netezza run-time processing

- overview [25](#)
- parameterization [26](#)
- partitioning [25](#)

Netezza write operation

- creating [24](#)

Netezza write operation properties

- advanced [20](#)
- ports [19](#)
- run-time [19](#)

P

PowerExchange for Netezza

- overview [7](#)

PowerExchange for Netezza configuration

- overview [8](#)
- prerequisites [8](#)