



Informatica® Cloud Data Integration

Netezza Connector

© Copyright Informatica LLC 2018, 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: 2024-04-10

Table of Contents

Preface	5
Informatica Resources.	5
Informatica Documentation.	5
Informatica Intelligent Cloud Services web site.	5
Informatica Intelligent Cloud Services Communities.	5
Informatica Intelligent Cloud Services Marketplace.	6
Data Integration connector documentation.	6
Informatica Knowledge Base.	6
Informatica Intelligent Cloud Services Trust Center.	6
Informatica Global Customer Support.	6
Chapter 1: Introduction to Netezza Connector	7
Netezza Connector assets.	7
Administration of Netezza Connector.	8
Download the Netezza JDBC driver.	8
Download the Netezza ODBC driver.	8
Chapter 2: Connections for Netezza	10
Netezza connection properties.	10
Chapter 3: Mappings and mapping tasks with Netezza Connector	12
Netezza sources in mappings.	12
Partitioning for Netezza sources.	14
Netezza targets in mappings.	14
Partitioning for Netezza targets.	16
Netezza target data update.	17
Update as insert.	17
Update else insert.	17
SQL statements.	18
Rules and guidelines for connection and mappings.	19
Chapter 4: Netezza ODBC SQL ELT optimization	20
Full SQL ELT optimization.	20
Source SQL ELT optimization.	21
SQL ELT functions.	21
Netezza SQL ELT transformations.	22
Configuring a Netezza ODBC connection for SQL ELT optimization.	23
Configuring the Netezza ODBC driver on Windows.	23
Configuring the Netezza ODBC driver on Linux.	26
Creating an ODBC connection.	26

Configuring SQL ELT optimization.	28
Rules and guidelines for Netezza ODBC connection	28
Chapter 5: Data type reference.	29
Netezza and transformation data types.	29
Index.	31

Preface

Use *Netezza Connector* to learn how to read from or write to Netezza by using Cloud Data Integration. Learn to create a Netezza connection, develop mappings, and run mapping tasks and dynamic mapping tasks in Cloud Data Integration. You can also learn how to configure SQL ELT optimization to a Netezza database using a Netezza ODBC connection.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Documentation

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 Netezza Connector

You can use Netezza Connector to read data from and write data of large volumes to a Netezza database.

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

Netezza Connector uses the Netezza ODBC driver to connect to the Netezza database. The Secure Agent 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 Secure Agent 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 Secure Agent 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 Secure Agent then writes the data from the staging table to the Netezza target.

You can switch mappings to advanced mode to include transformations and functions that enable advanced functionality.

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 Netezza Connector to read sales information from the Netezza database and use it for business analysis and decision making.

Netezza Connector assets

Create assets in Data Integration to integrate data using Netezza Connector.

When you use Netezza Connector, you can include the following Data Integration assets:

- Dynamic mapping task
- Mapping
- Mapping task

For more information about configuring assets and transformations, see *Mappings, Transformations, and Tasks* in the Data Integration documentation.

Administration of Netezza Connector

As a user, you can use Netezza Connector after the organization administrator ensures that users have access to the Secure Agent directory that contains the success and error files. This directory path must be the same on each Secure Agent machine in the run-time environment.

Before you can use Netezza Connector, you must also perform the following tasks:

- Install the Netezza client on Windows or Linux.
- Verify that you have the following privileges on the Netezza database:
 - CREATE TABLE
 - CREATE EXTERNAL TABLE
 - DELETE
 - DROP
 - INSERT
 - LIST
 - SELECT
 - TRUNCATE
 - UPDATE

Download the Netezza JDBC driver

1. Download the Netezza JDBC driver version from the IBM website:
To download the Netezza JDBC driver on Windows, follow the download instructions from the following Knowledge Base article: <https://kb.informatica.com/howto/6/Pages/23/619186.aspx>
If you want to use the Netezza JDBC driver on Linux, you can use the Netezza JDBC driver downloaded for Windows on the Linux machine.
2. After you download the Netezza JDBC driver, navigate to the following location on the Secure Agent location: `<Secure Agent installation directory>/apps/Data_Integration_Server/ext/`, and then manually create the following directory structure:
`deploy_to_main/bin/rdtm-extra/Netezza`
3. Copy the Netezza JDBC driver jar file, `nzjdbc.jar`, to the following directory you created on the Secure Agent machine:
`<Secure Agent installation directory>/apps/Data_Integration_Server/ext/
deploy_to_main/bin/rdtm-extra/Netezza`
4. Restart the Secure Agent.

Download the Netezza ODBC driver

1. Download the Netezza ODBC driver version from the IBM website:
To download the Netezza ODBC driver, follow the download instructions from the following Knowledge Base article:
https://knowledge.informatica.com/s/article/HOW-TO-Download-the-Netezza-ODBC-driver?language=en_US
2. Perform the following tasks to use the Netezza ODBC driver based on the operating system:
 - On Windows, verify if the NetezzaSQL driver appears in the ODBC Data Source Administrator driver list.

- On Linux, add the driver entries in the `odbcinst.ini` file in the Secure Agent installation directory:
The following code shows a sample entry:

```
[NetezzaSQL]
Driver          = /data/home/adputf_9/cloud_td/Netezza/installer/linux64/lib64/
libnzodbc.so
Setup          = /data/home/adputf_9/cloud_td/Netezza/installer/linux64/lib64/
libnzodbc.so
APILevel       = 1
ConnectFunctions = YYN
Description    = Netezza ODBC driver
DriverODBCVer  = 03.51
DebugLogging   = true
LogPath        = /tmp
UnicodeTranslationOption = utf8
CharacterTranslationOption = all
PreFetch       = 256
Socket         = 16384
```

CHAPTER 2

Connections for Netezza

Create a Netezza connection to securely read data from or write data to Netezza. You can use Netezza connections to specify sources and targets in mappings and mapping tasks.

Create a connection and associate it with a mapping or mapping task. Define the source and target properties to read data from or write data to Netezza.

You can create a Netezza connection on the **Connections** page and use it in the Mapping Designer when you create a mapping. The connection becomes available to the entire organization.

Netezza connection properties

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

The following table describes the Netezza connection properties:

Connection 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	Netezza.
Runtime Environment	The name of the runtime environment where you want to run tasks. Specify a Secure Agent, Hosted Agent, or serverless runtime environment.
Database	The name of the Netezza database.
Schemaname	The schema used for the Netezza source or target. Schema name is case sensitive.
Servername	The Netezza database host name.
Port	Network port number used to connect to the database server. Default is 1521.

Connection property	Description
Driver	The Netezza ODBC driver name, NetezzaSQL, used to connect to the Netezza database.
Runtime Additional Connection Configuration	Additional run-time attributes required to fetch data. For example, <code>securityLevel=preferredUnSecured;caCertFile =</code>
Metadata Additional Connection Configuration	The values to set the optional properties of the JDBC driver to fetch the metadata.
Username	Database user name with the appropriate read and write database permissions to access the database.
Password	Password for the database user name.

CHAPTER 3

Mappings and mapping tasks with Netezza Connector

When you create a mapping, you can configure a Source or Target transformation to represent a Netezza object.

In advanced mode, the Mapping Designer updates the mapping canvas to include transformations and functions that enable advanced functionality.

Netezza sources in mappings

In a mapping, you can configure a Source transformation to represent a Netezza source.

The following table describes the Netezza source properties that you can configure in a Source transformation:

Property	Description
Connection	Name of the source connection.
Source type	Type of the source object. Select Single Object, Multiple Objects, or Parameter.
Object	Name of the source object.

The following table describes the Netezza filter options that you can configure in a Source transformation:

Property	Description
Filter	<p>Filter value 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.</p> <p>You can specify the following filter conditions:</p> <ul style="list-style-type: none">- Not parameterized. Use a basic filter to specify the object, field, operator, and value to select specific records.- Completely parameterized. Use a parameter to represent the field mapping.- Advanced. Use an advanced filter to define a more complex filter condition on a dimension or metric.

The following table describes the Netezza advanced source properties that you can configure in a Source transformation:

Property	Description
Pipe Directory Path	<p>Path for the Secure Agent to create the pipe for the external table.</p> <p>If you do not specify the path, the Secure Agent uses the following directory to create the pipe for the external table:</p> <pre><Secure Agent installation directory>/apps/Data_Integration_Server/data/temp</pre>
Delimiter	<p>Required. Delimiter separates successive input fields.</p> <p>You can enter any value supported by the Netezza Performance Server. The value can be a part of the data for the Netezza source.</p> <p>Default is .</p>
NullValue	<p>NullValue parameter of an external table.</p> <p>The Secure Agent uses the NullValue internally. Maximum value is one character.</p> <p>Default is blank.</p>
EscapeCharacter	<p>Escape character of an external table.</p> <p>If the data contains NULL, CR, and LF characters in the Char or Varchar field, add an escape character in the source data before you extract. Enter an escape character before the data.</p> <p>The supported escape character is backslash (\).</p>
Socket Buffer Size	<p>Required. The socket buffer size that you must set to increase the mapping performance.</p> <p>Set from 25 to 50% of the DTM buffer size You might need to test different settings for optimal performance. Enter a value between 4096 and 2147483648 bytes.</p> <p>Default is 8388608 bytes.</p>
SQL Query Override	<p>Overrides the default query.</p> <p>Enter the SQL query that the Secure Agent must use to query data from the Netezza source.</p>
Owner Name	The Netezza schema name.
Source Table Name	<p>Overrides the source table name.</p> <p>For example, you can parameterize the source table name to override the table name in the mapping.</p>
PreSQL	Pre-SQL statement that the Secure Agent runs before reading data from the source.
PostSQL	Post-SQL command that the Secure Agent runs after reading data from the source.
Tracing Level	<p>Amount of detail that appears in the log for the transformation.</p> <p>Use the following tracing levels:</p> <ul style="list-style-type: none"> - Terse - Normal - Verbose Initialization - Verbose <p>Default is normal.</p>

Partitioning for Netezza sources

When you read data from Netezza sources, you can configure passthrough partitioning to read data in parallel and optimize the mapping performance at run time.

The Secure Agent distributes rows of data based on the number of partitions you define. Click the **Partitions** tab in the Source transformation to define the number of partitions you want to use to read data in parallel.

Netezza targets in mappings

To write data to Netezza, configure a Netezza object as the target in a mapping.

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

The following table describes the target properties that you can configure in a Target transformation:

Property	Description
Connection	Name of the target connection.
Target Type	Type of the target object. Select Single Object, Multiple Objects, or Parameter.
Object	Name of the target object.
Operation	Target operation. Select Insert, Update, Upsert, Delete, or Data Driven.

The following table describes the Netezza advanced target properties:

Property	Description
Pipe Directory Path	Path for the Secure Agent to create the pipe for the external table. If you do not specify the path, the Secure Agent uses the following directory to create the pipe for the external table: <Secure Agent installation directory>/apps/Data_Integration_Server/data/temp
Error Log Directory Name	Not applicable.
Enable Insert	The Secure Agent inserts rows into the target table and enforces key constraints. Default is selected.
Enable Delete	The Secure Agent deletes rows from the target table. Default is selected.

Property	Description
Enable Update	<p>Required. The Secure Agent updates the rows based on the update strategy option that you specify. Select one of the following values:</p> <ul style="list-style-type: none"> - Update as Update. The Secure Agent updates all rows flagged for update. - Update as Insert. The Secure Agent inserts all rows flagged for update. - Update else Insert. The Secure Agent first updates all rows flagged for update if they exist in the target. It then inserts the remaining rows marked for insert. - None. The Secure Agent does not update any row. <p>Default is None.</p> <p>When you use the Update property, the Secure Agent does not enforce key constraints and writes duplicate rows with the same primary key into the target table.</p>
Truncate Target Table Option	<p>The Secure Agent truncates the target before loading.</p> <p>Run the truncate table command.</p> <p>Default is disabled.</p> <p>You cannot use the Truncate Target Table option when you create a new target at runtime.</p>
Delimiter	<p>Required. Set the delimiter to any value supported by the Netezza Performance Server. The delimiter separates successive input fields. The value must not be a part of the input data.</p> <p>Default is .</p>
NullValue	<p>NullValue parameter of the external table. The Secure Agent uses the NullValue internally.</p> <p>Maximum value is one character.</p> <p>Default is blank.</p>
EscapeCharacter	<p>Escape character of the external table.</p> <p>If the data contains NULL, CR, and LF characters in the Char or Varchar field, you need to add an escape character for these fields before loading.</p> <p>Enter a backslash (\) as the escape character.</p>
Quoted Value	<p>QUOTEDVALUE parameter of the external table.</p> <p>Select SINGLE or DOUBLE to enclose the field in single or double quotes. Select NO to omit quotes.</p> <p>Default is NO.</p> <p>The quoted value is not a part of the data.</p>
Ignore Key Constraints	<p>Ignores constraints on primary key fields.</p> <p>When you select this option, the Secure Agent can write duplicate rows with the same primary key to the target.</p> <p>Default is disabled.</p> <p>The Secure Agent ignores this value when the target operation is "update as update" or "update else insert."</p>
Duplicate Row Handling Mechanism	<p>Determines how the Secure Agent handles duplicate rows. Select one of the following values:</p> <ul style="list-style-type: none"> - First Row. The Secure Agent passes the first row to the target and rejects the rows that follow with the same primary key. - Last Row. The Secure Agent passes the last duplicate row to the target and discards the rest of the rows. <p>Default is First Row.</p>
Bad File Name	Not applicable.

Property	Description
Socket Buffer Size	Required. The socket buffer size to set to increase session performance. Set a size of 25 to 50% of the DTM buffer size. You might need to test different settings for optimal performance. Enter a value between 4096 and 2147483648 bytes. Default is 8388608 bytes.
Control Character	Required. 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	Required. CRINSTRING parameter to transfer data containing carriage returns (CR). You can enter a non escape CR in Char or Varchar fields. To load the control characters present in the Char and Varchar fields, set the CTRLCHARS and CRINSTRING parameters to TRUE in the session properties for the Netezza source. Default is TRUE.
Table Name Prefix	Prefix that the Secure Agent 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.
PreSQL	Pre-SQL statement that the Secure Agent runs before writing data to the target.
PostSQL	Post-SQL statement that the Secure Agent runs after writing data to the target.
Success File Directory	Not applicable.
Error File Directory	Not applicable.
Forward Rejected Rows	Determines whether the transformation passes rejected rows to the next transformation or drops rejected rows. By default, the agent forwards rejected rows to the next transformation.

Partitioning for Netezza targets

When you use a Netezza connection in a mapping to write data to Netezza, you can configure partitioning for the Netezza target.

In the mapping, you do not have to specify the number of partitions for the target separately. The Secure Agent considers the same number of partitions for the target as the number defined for the source partitioning.

Netezza target data update

The Secure Agent updates target rows based on the update options and the duplicate row handling.

Update as insert

When you configure the mapping to update as insert rows, the Secure Agent uses the following process to update target rows:

- If the source key value matches a target key value, the Secure Agent does not insert the source row.
- If the source primary key value does not exist in the target, the Secure Agent inserts the source row.

The following table describes how the Secure Agent updates the target:

Source Data	Target Data	Updated Target Data	Comment
1,a,1a1	-	-	The source primary key is found in the target. The row is not inserted.
1,b,1b1	-	1,b,1b1	Inserts 1,b,1b1.
1,a,1a2	-	-	The source primary key is found in the target. The row is not inserted.
1,c,1c1	1,c,1c1	1,c,1c1	The source primary key is found in the target. The existing row 1,c,1c1 is retained. No insert is required.
1,d,1d1	-	1,d,1d1	Inserts 1,d,1d1.
1,a,1a3	1,a,1a3	1,a,1a3	The source primary key is found in the target. The existing row 1,a,1a3 is retained. No insert is required.

Note: In the pair of values, the first two values are the primary key, for example 1 (primary key), a (primary key), 1a1. The mapping is configured to consider key constraints.

Update else insert

When you configure the mapping to update else insert rows, the Secure Agent uses the following process to update target rows:

- If the source key value matches a target key value, the Secure Agent updates *each* target row. It updates with the first or last source row matched, based on how you configure duplicate row handling.
- If the source primary key value does not exist in the target, the Secure Agent inserts the source row.

The following table describes how the Secure Agent updates the target:

Source Data	Target Data	Updated Target Data	Comment
1,2	1,6	1,2	Updates 1,6 with 1,2. The source primary key is found in the target. The target row is updated based on duplicate row handling to use first row.
1,3	1,8	1,2	Updates 1,8 with 1,2. Duplicate row handling is configured to update with first source row. Subsequent target rows with primary key "1" are updated with first source row.
2,4	-	2,4	Inserts 2,4. The source primary key is not found in the target. The row is inserted.
2,5	-	-	Drops 2,5. The source primary key is found in the target, and first duplicate row has been updated in the target.
-	3,7	3,7	Retains 3,7. No update required.

Note: In the pair of values, the first value is the primary key, for example 1 (primary key), 2.

SQL statements

When you create a mapping in Cloud Data Integration, you can specify SQL statements in the advanced properties for a Netezza source or target.

You can define the following SQL statements for the source:

Pre-SQL

When you define a pre-SQL statement for a Netezza source in a mapping, the Secure Agent runs the pre-SQL statement before reading data from Netezza.

Post-SQL

When you define a post-SQL statement for a Netezza source in a mapping, the Secure Agent runs the post-SQL statement after reading data from Netezza.

You can define the following SQL statements for the target:

Pre-SQL

When you configure a pre-SQL statement for a Netezza target, the Secure Agent runs the pre-SQL statement before writing data to Netezza.

Post-SQL

When you configure a post-SQL statement for a Netezza target, the Secure Agent runs the post-SQL statement after writing data to the target.

Mapping Scenario with SQL Statements Defined for Netezza Sources and Targets

Consider a mapping where you want to read data from a Netezza source and write data to a Netezza target and you want to define SQL statements both for the Netezza source and target:

1. Define the following pre-SQL statement for the Netezza source: `delete from employee_tgt where job_id is null;`
2. Define the following post-SQL statement for the Netezza source: `update employee set first_name = 'First_name' where employee_id = 2`
3. Define the following pre-SQL statement for the Netezza target: `insert into emp_dept(emp_id,ename,dname) select employee_id,first_name,dept_name from department1,employee where employee.department_id = department1.dept_id`
4. Define the following post-SQL statement for the Netezza target: `insert into Staging_SalesData select * from SalesData where sales_dateTime > CURRENT_TIMESTAMP - 10`

When you run the mapping, the Secure Agent uses the following sequence for the tasks:

1. Runs the pre-SQL statement in the source Netezza database and deletes data of employees whose job IDs show null.
2. Runs the pre-SQL statement in the target Netezza database to update the first name of employees whose employee ID is 2.
3. Reads the data from the Netezza source database.
4. Writes the data to the Netezza target database.
5. Runs the post-SQL statement for staging data populated for another mapping in the Netezza source database and reads data from department1 and employee table based on the department ID.
6. Runs the post-SQL statement to insert additional sales data available from the last 10 days into the sales data staging table.

Rules and guidelines for connection and mappings

Consider the following rules and guidelines for Netezza connection and mappings:

- Schema name is case sensitive.
- You cannot use the Truncate Target Table option when you create a new target at runtime.

CHAPTER 4

Netezza ODBC SQL ELT optimization

When you read data from a Netezza source, transform the data, and write the data to a target, you can configure SQL ELT optimization to push the transformation logic to the source or target database system.

You can configure SQL ELT optimization when you use a Netezza ODBC connection that uses Netezza ODBC drivers in the mapping.

If the source and target databases are the same, you can configure full SQL ELT optimization for improved performance.

When the Secure Agent applies SQL ELT optimization, it pushes transformation logic to a database. The Secure Agent translates the transformation logic into SQL queries and sends the SQL queries to the database. The database runs the SQL queries to process the transformations.

SQL ELT optimization improves mapping performance when the database can process the transformation logic faster than the Secure Agent. The Secure Agent also reads less data from the database.

The amount of transformation logic that the Secure Agent pushes to the database depends on the database, the transformation logic, and the mapping task. The Secure Agent processes all transformation logic that it cannot push to a database.

When you configure SQL ELT optimization for the mapping, the Secure Agent analyzes the optimized mapping from the source to the target or until it reaches a downstream transformation that it cannot push to the source database. The Secure Agent generates and executes a SELECT statement for each source that has transformation logic pushed down. Then, it reads the results of this SQL query and processes the remaining transformations in the mapping.

You cannot use SQL ELT optimization when you run a mapping in advanced mode.

Full SQL ELT optimization

The ODBC connection supports Source SQL ELT optimization when you use the Netezza ODBC driver in a mapping.

You can configure full SQL ELT optimization only when the source and target are in the same database.

When the Secure Agent applies full SQL ELT optimization, it pushes all the transformation logic in the mapping to the target database. You can configure full SQL ELT optimization in the SQL ELT Optimization section on the **Runtime Options** tab of the mapping task.

Full SQL ELT optimization is ideal when the source and target are in the same connections. For example, if a mapping contains a Netezza source and a Netezza target, configure full SQL ELT optimization to push all the transformation logic for processing from a Netezza source database to a Netezza target database.

Source SQL ELT optimization

The ODBC connection supports Source SQL ELT optimization when you use the Netezza ODBC driver in a mapping.

When the Secure Agent applies source SQL ELT optimization, it analyzes the mapping from source to target or until it reaches a downstream transformation it cannot push to the source database.

The Secure Agent generates and executes a SELECT statement based on the transformation logic for each transformation it can push to the database. Then, it reads the results of this SQL query and processes the remaining transformations.

You can configure a mapping to use source SQL ELT optimization if the source and target reside in different databases. For example, if a mapping contains a Netezza source and a Netezza target, you can configure source SQL ELT optimization to push some transformation logic for processing to the Netezza source.

You can configure source SQL ELT optimization in the SQL ELT Optimization section on the **Runtime Options** tab of the mapping task.

SQL ELT functions

The following table summarizes the availability of SQL ELT functions in a Netezza source or target:

Functions	SQL ELT optimization Type
ABS()	Source, Full
AVG()	Source, Full
COS()	Source, Full
COUNT()	Source, Full
DATE_COMPARE()	Source, Full
DECODE()	Source, Full
EXP()	Source, Full
IIF()	Source, Full
IN()	Source, Full
ISNULL()	Source, Full

Functions	SQL ELT optimization Type
LOWER()	Source, Full
MAX()	Source, Full
MIN()	Source, Full
POWER()	Source, Full
SIN()	Source, Full
SQRT()	Source, Full
SUM()	Source, Full
TAN()	Source, Full
TO_DECIMAL()	Source, Full
TO_FLOAT()	Source, Full
TO_INTEGER()	Source, Full
UPPER()	Source, Full

Netezza SQL ELT transformations

When you configure SQL ELT optimization, the Secure Agent tries to push each transformation to the database.

The Secure Agent can push the following transformation logic to a Netezza source or target:

Transformations	Supported SQL ELT Type
Aggregator	Source, Full
Expression	Source, Full
Filter	Source, Full
Joiner	Source, Full
Sorter	Source, Full
Union	Source, Full
Router	Full

Configuring a Netezza ODBC connection for SQL ELT optimization

You can set the SQL ELT optimization for the ODBC connection that uses the Netezza ODBC driver to enhance the mapping performance.

You must first configure the Netezza ODBC drivers and create a Netezza ODBC data source name. Netezza supports Netezza ODBC drivers on Windows and Linux systems.

After you configure the Netezza ODBC DSN, you must create an ODBC connection to connect to Netezza. In the ODBC connection properties, select **ODBC Subtype** as `Other` and the **Driver Manager** as `Data Direct`.

For more information about creating an ODBC connection to connect to Netezza, see the "Creating an ODBC Connection" topic in the help for ODBC Connector.

Use the configured Netezza ODBC connection in a mapping, and select the value of the SQL ELT Optimization property as **Full** or **To Source** in the SQL ELT optimization of the mapping task.

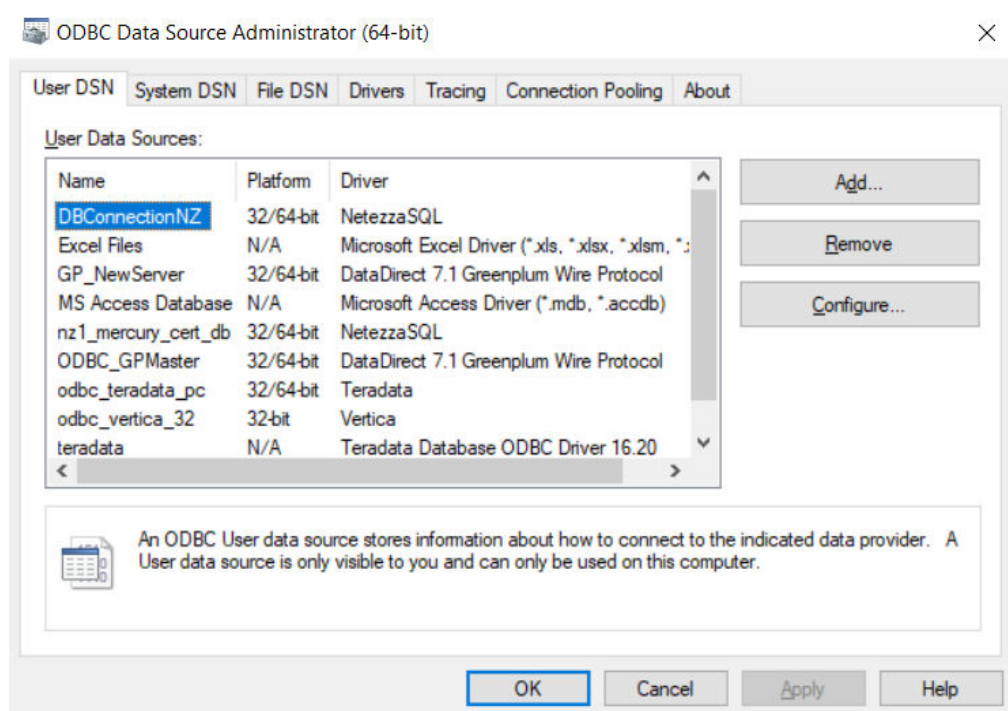
Configuring the Netezza ODBC driver on Windows

Before you establish an ODBC connection to connect to Netezza on Windows, you must configure the Netezza ODBC driver and create the Netezza data source name.

Perform the following steps to configure an ODBC connection on Windows:

1. Open the folder in which the ODBC data source file is installed.
2. Run the `odbcad32.exe` file.
The **ODBC Data Source Administrator** dialog box appears.
3. Click **System DSN**.
The **System DSN** tab appears.

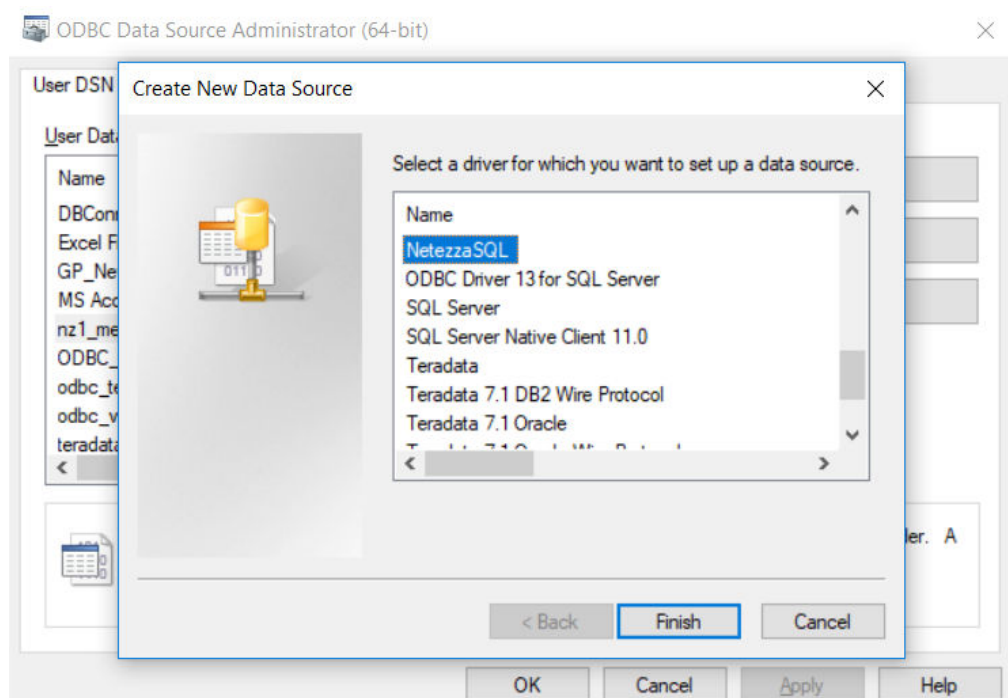
The following image shows the System DSN tab on the ODBC Data Source Administrator (64-bit) dialog box:



4. Click **Add**.

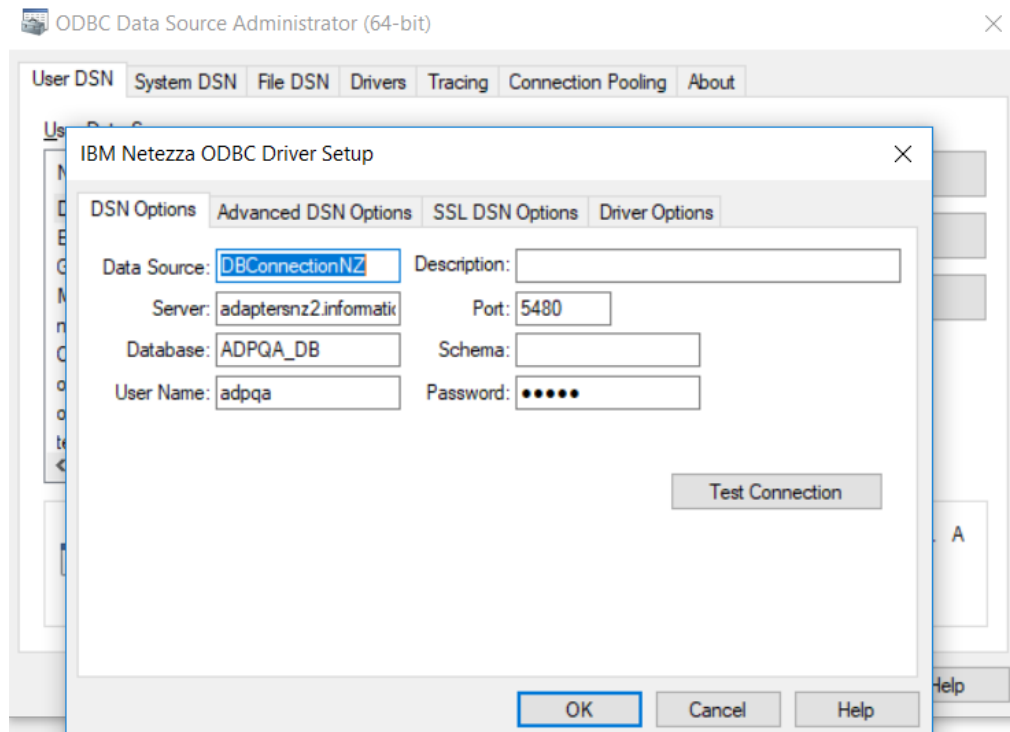
The **Create New Data Source** dialog appears.

The following image shows the Create New Data Source dialog box where you can select the Netezza data source:



5. Select **NetezzaSQL** and click **Finish**.
6. Click **Configure**.

The Netezza Configuration dialog box appears.
The following image shows the DSN options tab:



7. Specify the following DSN properties:

Property	Description
Data Source	Name of the data source.
Description	Description for the data source.
Server	Domain name of your Netezza account.
Port	Port number used to connect to Netezza.
Database	Name of the Netezza database.
Schema	Name of the Netezza schema.
User Name	Username to access the Netezza database.
Password	Password to access the Netezza database.

8. Click **OK**.

The Netezza ODBC driver is configured successfully on Windows.

Configuring the Netezza ODBC driver on Linux

Before you can run tasks to connect to Netezza using the ODBC connection from Linux, you must set the ODBCINI, ODBCINST, and LD_LIBRARY_PATH environmental variables for the 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/cloud_td/ODBCINI/odbc.ini"
```
2. To set the ODBCINST environment variable, use the following syntax:

```
setenv ODBCINST /data/home/qamercury/cloud_td/ODBCINI/odbcinst.ini
```
3. To set the LD_LIBRARY_PATH environment variable, use the following syntax:

```
setenv LD_LIBRARY_PATH ".:./export/qa_adp/thirdparty/netezza/linux.64/  
lib64:$LD_LIBRARY_PATH"
```
4. To set the NZ_ODBC_INI_PATH environment variable, use the following syntax:

```
setenv NZ_ODBC_INI_PATH "/path/to/directory/where/we/have/odbc.ini/file"
```
5. Add entries for the Netezza data sources in the `odbc.ini` file.

The following section shows a sample entry in the `odbc.ini` file:

```
[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
```

6. Restart the Secure Agent after you configure the environment variables.

Creating an ODBC connection

You must create an ODBC connection to connect to Netezza after you configure the Netezza ODBC drivers.

Perform the following steps to create a Netezza ODBC connection on the **Connections** page:

1. In Administrator, click **Connections**.
The Connections page appears.
2. Click **New Connection**.
The **New Connection** page appears.

The following image shows the New Connection page:

3. Configure the following connection details in the **Connection Details** section:

Property	Description
Connection Name	Name of the ODBC connection. For example, Netezza_ODBC.
Description	Description of the connection.
Type	Type of the connection. Select the type of the connection as ODBC .

4. Configure the following connection details in the **ODBC Connection Properties** section:

Property	Description
Runtime Environment	Runtime environment that contains the Secure Agent you can use to access the system.
User Name	Username to log in to the Netezza database.

Property	Description
Password	Password to log in to the Netezza database.
Data Source Name	Enter the name of the ODBC data source name that you created for the Netezza database.
Schema	Name of the Netezza schema.
Code Page	The code page of the database server defined in the connection.
ODBC Subtype	Enter the value of the ODBC Subtype field as Other .
Driver Manager for Linux	The driver that the Netezza ODBC driver manager sends database calls to.

The Netezza ODBC connection is created successfully.

Configuring SQL ELT optimization

In a Netezza mapping task, use the configured Netezza ODBC connection.

Perform the following steps to configure SQL ELT optimization for Netezza sources or targets:

1. In the **Runtime Options** tab of the Mapping task, navigate to the SQL ELT Optimization section.
2. From the SQL ELT Optimization list, select the required type of SQL ELT optimization.

Rules and guidelines for Netezza ODBC connection

Consider the following rules and guidelines for SQL ELT optimization to a Netezza database:

- When you configure SQL ELT optimization in a mapping to write data that contains numeric fields to a Netezza target, the Secure Agent rounds large-range numeric data (38,10) to the first decimal place.
- You cannot push down data of the Timestamp data type to a Netezza target database.

CHAPTER 5

Data type reference

Data Integration uses the following data types in synchronization tasks, mappings, and mapping tasks with Netezza:

Netezza Native Data Types

Netezza 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 runtime environment uses to move data across platforms. Transformation data types appear in all transformations in synchronization tasks, mappings, and mapping tasks.

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.

Netezza and transformation data types

The following table lists the Netezza data types that the runtime environment supports and the corresponding transformation data types:

Netezza Datatype	Range	Transformation Datatype	Range
BigInt	Precision 19, scale 0	Bigint	From -9,223,372,036,854,775,808 through 9,223,372,036,854,775,807 Precision of 19, scale of 0 Integer value
Bool	True or false, on or off, 0 or 1, yes or no	String	Precision 1
ByteInt	Precision 3, scale 0	Small Integer	Precision 5, scale 0
Char	Single character	String	From 1 through 104,857,600 characters
Date	ANSI SQL date	Date/Time	Jan. 1, 0001 A.D. to Dec. 31, 9999 A.D. (precision to the nanosecond)
Float8	Precision 15	Double	Precision 15

Netezza Datatype	Range	Transformation Datatype	Range
Float4	Precision 6, scale 0	Double	Precision 15
Integer	Precision 10, scale 0	Integer	Precision 10, scale 0
NChar(m)	Single character Used for storing UTF-8 data	nString	From 1 through 104,857,600 characters
NVarchar(m)	BVarchar (length) Non-blank-padded string, variable storage length Used for storing UTF-8 data	nString	From 1 through 104,857,600 characters
Numeric	Numeric (precision, decimal), arbitrary precision number Precision must be between 1 and 38	Decimal	Precision from 1 through 28 digits, scale from 0 through 28
Real	Precision 6, scale 0	Real	Precision of 7, scale of 0 Double-precision floating-point numeric value
SmallInt	Precision 5, scale 0	Small Integer	Precision 5, scale 0
Time	hh:mm:ss. ANSI SQL time	Date/Time	Jan. 1, 0001 A.D. to Dec. 31, 9999 A.D. (precision to the microsecond)
Timestamp	Precision 26, scale 6	Date/Time	Jan. 1, 0001 A.D. to Dec. 31, 9999 A.D. (precision to the microsecond)
Varchar	Varchar (length) Non-blank-padded string, variable storage length	String	From 1 through 104,857,600 characters

INDEX

C

Cloud Application Integration community
URL [5](#)
Cloud Developer community
URL [5](#)
Configuring Netezza ODBC Connection
Windows [23](#)
configuring Netezza ODBC driver
Linux [26](#)
connections
Netezza [10](#)

D

Data Integration community
URL [5](#)
data type reference
overview [29](#)

I

Informatica Global Customer Support
contact information [6](#)
Informatica Intelligent Cloud Services
web site [5](#)

K

key constraints
example [17](#)

M

maintenance outages [6](#)

N

Netezza
connection properties [10](#)
supported object types [7](#)
supported task types [7](#)
Netezza connections
administration [8](#)
Netezza Connections
overview [10](#)
Netezza Connector
overview [7](#)

Netezza Data Types
data types [29](#)
Netezza objects
mapping [12](#)
Netezza ODBC connection
configuration [23](#)
Netezza sources
mapping [12](#)
Netezza targets
mappings [14](#)

S

SQL ELT method
full SQL ELT optimization [20](#)
SQL ELT optimization
create an ODBC connection [26](#)
functions [21](#)
Other subtype [26](#)
SQL ELT Optimization
Configuration [28](#)
SQL ELT optimization method
source pushdown [21](#)
status
Informatica Intelligent Cloud Services [6](#)
system status [6](#)

T

trust site
description [6](#)

U

update as insert
description for Netezza [17](#)
update else insert
description for Netezza [17](#)
update strategy
example [17](#)
update as insert for Netezza [17](#)
update else insert for Netezza [17](#)
upgrade notifications [6](#)

W

web site [5](#)