



TIBCO® Data Virtualization

Netezza Adapter Guide

Version 8.7.0 | October 2023

Contents

Contents	2
TDV Netezza Adapter	4
Introduction	4
Obtain and Install the Driver for Netezza	4
Datasource Configuration	5
Basic Tab	5
Advanced Tab	9
Netezza Setup	22
Netezza Data Source Limitations	23
Native (Bulk) Caching DDL Creation Limitations	23
Netezza Caching Tips from an Expert	24
Data Type Mappings	24
Netezza to TDV Data Types	24
Netezza Cache Mapping	26
Netezza Function Support	28
Netezza Aggregate Function Support	28
Netezza Analytic Aggregate Function Support	30
Netezza Analytic Function Support	31
Netezza Binary Function Support	32
Netezza Character Function Support	33
Netezza Conditional Function Support	35
Netezza Conversion Function Support	35
Netezza Date Function Support	36
Netezza Numeric Function Support	37
Netezza Phonetic Function Support	40
Netezza Statistical Analytic Aggregate Function Support	41
Netezza Time Function Support	41

References	42
TIBCO Product Documentation and Support Services	44
How to Access TIBCO Documentation	44
How to Contact TIBCO Support	45
Release Version Support	45
How to Join TIBCO Community	46
Legal and Third-Party Notices	47

TDV Netezza Adapter

Introduction

This section explains the various connection and configuration options of the datasource Netezza as well as the capabilities:

[Obtain and Install the Driver for Netezza](#)

[Datasource Configuration](#)

[Netezza Data Source Limitations](#)

[Data Type Mappings](#)

[Netezza Function Support](#)

[References](#)

Obtain and Install the Driver for Netezza

Obtain and install the JDBC driver for Netezza

Obtain the following JDBC driver for Netezza from the NPS system, or contact the support group at Netezza for the driver.

`nzjdbc<v>.jar`

1. Copy it to the TDV installation directory:

```
<TDV_install_dir>\conf\adapters\system\netezza_x_x
```

The `x_x` represents the version number of the Netezza JDBC driver being installed.

2. Restart the TDV Server.

Refer to the Installation guide for more details about the supported versions.

Datasource Configuration

This section explains the connection properties that are defined while setting up a datasource.

Basic Tab

The following table and the sections below lists and explains the connection properties that are in the Basic Tab of the New Data Source Window.

Datasource Name	The name of the Datasource.
Host	Name of the host machine or the host machine's IP address.
Port	Port number for the data source to connect with the host. The Port number for Netezza is 5480.
Database Name	Name or alias of the underlying data source. TDV Server uses this name to find and connect to the data source.
Login/User, Password	User name and password required to access the data source.
Pass-through Login	Flag to indicate whether pass-through login is enabled or not.
Transaction Isolation	The degree to which transactions are isolated from data modifications made by other transactions.

Datasource Name

The name of the data source.

Data Type

string

Default Value

""

Host

Name of the host machine or the host machine's IP address.

Data Type

string

Default Value

""

Port

The Port number

Data Type

string

Default Value

""

Database Name

Name or alias of the underlying data source. TDV Server uses this name to find and connect to the data source.

Data Type

string

Default Value

""

Login/User, Password

User name and password required to access the data source.

Data Type

string

Default Value

""

Remarks

When the data source is used as a target for cache tables or data ship, the user must also have permission to create tables, execute DDL, and perform other required tasks. Refer to the individual data source descriptions for details.

Transaction Isolation

The degree to which transactions are isolated from data modifications made by other transactions.

Data Type

string

Default Value

NONE

Remarks

Valid values are:

- Read Committed (default)—Nonrepeatable reads and phantom reads can occur.
- Serializable—Dirty reads, nonrepeatable reads, and phantom reads are prevented.
- None

Pass-through Login

Flag to indicate whether pass-through login is enabled or not.

Data Type

string

Default Value

""

Remarks

Disabled (default)—This allows automated provisioning of a connection pool. Open connection threads can be used by authorized users after the validation query verifies connection status. If pass-through login is disabled, the Save Password check box is not available.

Enabled—A new connection to the data source uses the credentials supplied by the client when data is requested from that data source for the first time. Subsequent requests by the same user reuse the existing connection. When another user attempts to connect to a data source, a new connection is created.

See “Managing Security for TDV Resources” in the TDV Administration Guide for details.

Authentication Type

Indicates the type of authentication used by the data source.

Data Type

String

Default Value

BASIC

Remarks

Select BASIC or Kerberos authentication method, where offered.

See the *TDV Administration Guide* for more information about Kerberos authentication.

Advanced Tab

The following table and the sections below lists and explains the connection properties that are in the Advanced Tab of the New Data Source Window.

Connection URL Pattern	A template for generating a URL to connect to the physical data source.
Connection URL String	The URL string generated from the connection URL pattern with the connection information you provide.
JDBC Connection Properties	Lets you specify property-value pairs to pass to the JDBC data source
Connection Pool Maximum Size	Maximum number of connections (both active and idle) allowed for the data source. When the maximum is reached, new requests must wait until a connection is available.
Connection Pool Minimum Size	Minimum number of connections in the pool even when the pool is inactive.
Connection Pool Idle Timeout	Number of seconds that a connection can remain idle without being dropped from the pool when there are

	more than the minimum number of connections.
Maximum Connection Lifetime	The number of minutes that a connection that was returned to the pool persists if there are more open connections than the minimum pool size.
Connection Validation Query	A data-source-specific query that the TDV query engine sends to see if the data source connection is valid. This query is executed every time a connection is checked out from the pool. Enter a query that returns quickly.
Execution Timeout	The number of seconds an execution query on the data source can run before being canceled.
Execute SELECTs Independently	Lets a SELECT statement be executed using a new connection from the connection pool, and committed immediately after completion. INSERT, UPDATE, and DELETE statements are executed using the same connection as part of the transaction.
Connection Checkout Procedure	A procedure that returns a valid SQL statement that can be used to initialize the connection.
Supports Star Schema	Check only if this data source supports very large predicates and very large cardinalities for star schema semijoins.
Max Source Side Cardinality for Semi Join	See the documentation for semijoins and the TDV Administration Guide for more information.
Min Target to Source Ratio for Semi Join	Sets the minimum target-to-source ratio of cardinality for semijoins. Refer to the TDV Administration Guide for more information.
Max Source Side of Semi Join To Use OR Syntax	See the documentation for semijoins and the TDV Administration Guide for more information.
Enable Native Data Loading	Lets the data source use its proprietary functionality to optimize performance.

Collation Sensitive	TDV does not use the SORT MERGE join algorithm if any data source involved in the join is marked Collation Sensitive.
Concurrent Request Limit	Works with the Massively Parallel Processing engine configuration parameters to control the amount of parallelization for the queries for a particular data source.
Is dataship target	Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source. Note: All Netezza data sources should be configured to act as data ship targets.
Lower bound/Upper bound for dataship	TDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.
Schema Path For Temp Tables	A relative path to set the location of the temp tables on the data source. It is the name of a schema in the data source.
Temp Table Prefix	A character string addition to temporary table names so that they are recognized if they are needed.

Connection URL Pattern

A template for generating a URL to connect to the physical data source.

Data Type

string

Default Value

jdbc:<DATA_SOURCE>//<HOST>:<PORT>/<DATABASE_NAME>

Remarks

TDV does not validate modifications at the time of configuration. The data source adapter might not validate changes.

Connection URL String

The URL string generated from the connection URL pattern with the connection information you provide.

Data Type

string

Default Value

“”

Remarks

This string is used by the JDBC adapter to connect to the physical data source. This field cannot be edited. For details, see the section “Connecting through JDBC Adapters” in the *TDV Administration Guide*.

JDBC Connection Properties

Lets you specify property-value pairs to pass to the JDBC data source.

Data Type

string

Default Value

“”

Remarks

Click to add custom connection properties for any JDBC data source. Commonly used properties are populated with default values. Use the Add Argument button to specify other properties and values.

TDV does not validate property names. Some data source adapters ignore invalid property names or values; others return an error.

The driver properties specify connection timeout settings required by specific drivers. To avoid leaving connections open indefinitely, specify properties explicitly for your data source.

Connection Pool Maximum Size

Maximum number of connections (both active and idle) allowed for the data source. When the maximum is reached, new requests must wait until a connection is available.

Data Type

Numeric

Default Value

100

Remarks

If the maximum number of connections is in use when a request comes in (even with pass-through authentication), the new request is blocked and queued until a connection is available or the Connection Pool Idle Timeout is reached.

If no connection was made available within the specified timeout, a check is made for an available connection by the same user. If none is available, the least recently used connection for another user is dropped and a new connection is opened.

Studio reuses pooled connections if they continue to be valid after changes (such as connection name), but JDBC requests are forced to use new connections if any part of the data source connection configuration has changed.

Connection Pool Minimum Size

Minimum number of connections in the pool even when the pool is inactive.

Data Type

Numeric

Default Value

0

Remarks

When a connection has been idle, a validation query is used to verify whether an open connection is still valid just prior to submission of a request. If the connection is invalid, the connection is discarded and another is used.

Connection Pool Idle Timeout

Number of seconds that a connection can remain idle without being dropped from the pool when there are more than the minimum number of connections.

Data Type

Numeric

Default Value

30

Maximum Connection Lifetime

The number of minutes that a connection that was returned to the pool persists if there are more open connections than the minimum pool size.

Data Type

Numeric

Default Value

30

Remarks

The duration is calculated from connection creation. Default value is 60 minutes. Set a smaller value if the pool is likely to run out of connections. Be sure to add a validation query. Set a larger value if you want the connections to be held for a longer period. Set a value of 0 to keep connections alive indefinitely.

Connection Validation Query

A data-source-specific query that the TDV query engine sends to see if the data source connection is valid. This query is executed every time a connection is checked out from the pool. Enter a query that returns quickly.

Data Type

string

Default Value

""

Remarks

If this query returns a non-error result, the data source connection is considered valid. If this query fails, the connection is discarded and a new connection is checked out from the available pool.

No one SELECT statement works with all data sources. To verify that TDV is running and that it can connect to the data source, devise a query against a published table from that data source.

Enable Native Data Loading

Lets the data source use its proprietary functionality to optimize performance.

Data Type

Bool

Default Value

True

Remarks

See the User Guide, Chapter About Data Source Native Load Performance Options” for more details,.

Collation Sensitive

TDV does not use the SORT MERGE join algorithm if any data source involved in the join is marked Collation Sensitive.

Data Type

Bool

Default Value

False

Remarks

None

Concurrent Request Limit

Works with the Massively Parallel Processing engine configuration parameters to control the amount of parallelization for the queries for a particular data source.

Data Type

Numeric

Default Value

0

Remarks

None

Is dataship target

Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.

Note: All Netezza data sources should be configured to act as data ship targets.

Data type

Bool

Default Value

False

Remarks

None

Lower bound/Upper bound for dataship

TDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.

Data Type

Numeric

Default Value

Lower bound - 50000

Upper bound - 5000000

Remarks

None

Schema Path For Temp Tables

A relative path to set the location of the temp tables on the data source. It is the name of a schema in the data source.

Data Type

String

Default Value

""

Remarks

This field is available only if you enable Is dataship target field.

Temp Table Prefix

A character string addition to temporary table names so that they are recognized if they are needed.

Data Type

String

Default Value

T

Remarks

This field is available only if you enable Is dataship target field.

Execution Timeout

The number of seconds an execution query on the data source can run before being canceled.

Data Type

Numeric

Default Value

0

Remarks

None

Execute SELECTs Independently

Lets a SELECT statement be executed using a new connection from the connection pool, and committed immediately after completion. INSERT and UPDATE statements are executed using the same connection as part of the transaction.

Data Type

Bool

Default Value

True

Remarks

None

Connection Checkout Procedure

A procedure that returns a valid SQL statement that can be used to initialize the connection.

Data Type

string

Default Value

""

Remarks

The signature of the initialization procedure should be:

```
(IN ds_name VARCHAR, OUT sqlText VARCHAR)
```

Give the full path to the procedure in the Connection Check-out Procedure box.

Max Source Side Cardinality for Semi Join

See the documentation for semijoins and the TDV Administration Guide for more information.

Data Type

Numeric

Default Value

""

Remarks

None

Max Source Side of Semi Join To Use OR Syntax

See the documentation for semijoins and the TDV Administration Guide for more information.

Data Type

Numeric

Default Value

2147483647

Remarks

None

Min Target to Source Ratio for Semi Join

Sets a minimum ratio to trigger use of semi join optimization.

Data Type

Numeric

Default Value

""

Remarks

None

Supports Star Schema

Check only if this data source supports very large predicates and very large cardinalities for star schema semijoins.

Data Type

Bool

Default Value

False

Remarks

Refer to the section Star Schema Semijoin in the User Guide, for more information.

Netezza Setup

The TDV Server can push Netezza SQL analytic functions, aggregate functions, regular expressions, and associated keywords.

To introspect data, TDV requires SELECT permissions on the following Netezza tables and views.

System Tables and Views	<CurrentCatalog>.<CurrentSchema> Tables and Views
<ul style="list-style-type: none"> ADMIN._T_AGGREGATE ADMIN._T_PROC 	<ul style="list-style-type: none"> _v_function _v_aggregate _V_JDBC_PROCEDURE_COLUMNS2 _V_JDBC_PKFK2 _V_JDBC_PROCEDURES2 _V_JDBC_PRIMARYKEYS2 _V_JDBC_INDEXINFO2 _V_JDBC_COLUMNS2 _V_JDBC_TABLES2 _V_JDBC_PROCEDURES2

After successful connection with the data source, Netezza tables, views, user-defined aggregates (UDAs), and user-defined functions (UDFs) are displayed for TDV introspection.

Netezza Data Source Limitations

- Placeholder parameters (designated by question marks) in prepared statements cannot be evaluated for use with the Netezza data ship optimization, because the variable must be resolved prior to submission.
- Federated queries with database-specific functions must be able to push that SQL directly to the data source, or the query will fail.
- Netezza UDA and UDF need to be assigned full resource names so they can be used in queries.
- Composite views can invoke Netezza UDA or UDF using full TDV resource names of the following general form:
`/users/composite/admin/Netezza/Aggregate/Builtin/"COUNT"()`
- Each query (including the initial costing of the query) is executed using its own run-time connection thread, which is returned to the pool when the query is completed.

Native (Bulk) Caching DDL Creation Limitations

When the data source tries to create a table that contains certain data types, some data types are not supported for particular data source and cache target combinations.

Data Source	Cache Target	Data Type Not Supported
DB2	Netezza 6	BLOB, CLOB
Sybase ASE 15	Netezza 6	TIMESTAMP, BINARY, IMAGE, TEXT, VARBINARY
SQL Server 2008	Netezza 6	VARBINARY, TIMESTAMP, BINARY, TEXT, IMAGE, NTEXT
Oracle	Netezza 6	BLOB, CLOB, NCLOB, LONG, LONGRAW

Data Source	Cache Target	Data Type Not Supported
Any data source	Vertica	Any data type (BINARY, CHAR, VARCHAR, BLOB, and so on) with length greater than 65000

Netezza Caching Tips from an Expert

When Netezza is the target of your cache, the experts recommend that you configure the Netezza database Transaction Isolation levels to the default value of Serializable.

Data Type Mappings

Netezza to TDV Data Types

Netezza's BIT data type is equivalent to a BOOLEAN type. However, it is not accepted in mathematical operations.

The following table shows the mapping from Netezza data types to TDV data types.

Netezza Data Type	TDV Data Type
BIGINT	BIGINT
BOOL or BOOLEAN	BOOLEAN (See Mapping of Native to TDV Data Types Across TDV Versions, page 555.)
BYTEINT	TINYINT
CHAR	CHAR

Netezza Data Type	TDV Data Type
DATE	DATE
DECIMAL	DECIMAL
DOUBLE PRECISION	DOUBLE
FLOAT	FLOAT
INT	INTEGER
INT1	TINYINT
INT2	SMALLINT
INT4	INTEGER
INT8	BIGINT
INTEGER	INTEGER
INTERVAL	VARCHAR
NCHAR	CHAR
NVARCHAR	VARCHAR
NUMERIC	DECIMAL
REAL	REAL
SMALLINT	SMALLINT
TIME	TIME
TIMETZ	TIMESTAMP VARCHAR [not in 5.0, 6.0]

Netezza Data Type	TDV Data Type
TIME WITH TIME ZONE	VARCHAR
TIMESTAMP	TIMESTAMP
VARCHAR	VARCHAR

Netezza Cache Mapping

This section discusses the data type mappings for caches stored on Netezza.

Netezza data types have these characteristics:

- IN predicate (with multiple values) with a subquery. It does not support multiple values on the right hand side. For example, Pushable (x,y) IN (select a,b from foo) :
Not Pushable - will be processed within TDV - (k, j) IN ((a, b), (c, d), (e, f))

Overrides are indicated in square brackets.

Data Type	Native Type
BIGINT	BIGINT, INT8 [not 5.0, 6.0]
BIT	BOOLEAN
BOOL	BOOLEAN
BOOLEAN	BOOLEAN
CHAR	CHAR(n), CHAR [not 5.0, 6.0]
DATE	DATE
DECIMAL	NUMERIC(p,s), DECIMAL [not 5.0, 6.0]
DOUBLE	DOUBLE PRECISION
FLOAT	FLOAT

Data Type	Native Type
INTEGER	INT, INT4, INTEGER
INTERVAL	VARCHAR
INTERVAL DAY	VARCHAR(30)
INTERVAL DAY TO HOUR	VARCHAR(30)
INTERVAL DAY TO MINUTE	VARCHAR(30)
INTERVAL DAY TO SECOND	VARCHAR(30)
INTERVAL HOUR	VARCHAR(30)
INTERVAL HOUR TO MINUTE	VARCHAR(30)
INTERVAL HOUR TO SECOND	VARCHAR(30)
INTERVAL MINUTE	VARCHAR(30)
INTERVAL MINUTE TO SECOND	VARCHAR(30)
INTERVAL MONTH	VARCHAR(9)
INTERVAL SECOND	VARCHAR(30)
INTERVAL YEAR	VARCHAR(9)
INTERVAL YEAR TO MONTH	VARCHAR(12)
NUMERIC	NUMERIC(p,s)
REAL	REAL
SMALLINT	SMALLINT, INT2
TIME	TIME, TIMEZ

Data Type	Native Type
TIMESTAMP	TIMESTAMP
TIMETZ(n)	TIMETZ [5.0, 6.0]
TINYINT	SMALLINT, INT1 [not 5.0, 6.0]
VARCHAR	VARCHAR(n)

Netezza Function Support

TDV supports the following types of functions for Netezza:

- [Netezza Aggregate Function Support](#)
- [Netezza Analytic Function Support](#)
- [Netezza Analytic Aggregate Function Support](#)
- [Netezza Binary Function Support](#)
- [Netezza Character Function Support](#)
- [Netezza Conditional Function Support](#)
- [Netezza Conversion Function Support](#)
- [Netezza Date Function Support](#)
- [Netezza Numeric Function Support](#)
- [Netezza Phonetic Function Support](#)
- [Netezza Statistical Analytic Aggregate Function Support](#)
- [Netezza Time Function Support](#)

Netezza Aggregate Function Support

TDV supports the aggregate functions listed in the table below for Netezza.

Netezza Aggregate Function	Notes
AVG	<p>Not supported:</p> <ul style="list-style-type: none"> • AVG (BOOLEAN) • AVG (BOOLEAN) DISTINCT • AVG (NCHAR) • AVG (NCHAR) DISTINCT • AVG (NVARCHAR) • AVG (NVARCHAR) DISTINCT
COUNT	
MAX	<p>Not supported:</p> <ul style="list-style-type: none"> • MAX (BOOLEAN) • MAX (BOOLEAN) DISTINCT • MAX (NCHAR) • MAX (NCHAR) DISTINCT • MAX (NVARCHAR) • MAX (NVARCHAR) DISTINCT <p>In version 6.0: analytic, with the same arguments not supported. MAX(NULL) analytic is NULL.</p>
MIN	<p>Not supported:</p> <ul style="list-style-type: none"> • MIN (BOOLEAN) • MIN (BOOLEAN) DISTINCT • MIN (NCHAR) • MIN (NCHAR) DISTINCT • MIN (NVARCHAR) • MIN (NVARCHAR) DISTINCT

Netezza Aggregate Function	Notes
SUM	<p>Not supported:</p> <ul style="list-style-type: none"> • SUM (BOOLEAN) • SUM (BOOLEAN).DISTINCT • SUM (NCHAR) • SUM (NCHAR).DISTINCT • SUM (NVARCHAR) • SUM (NVARCHAR).DISTINCT

Netezza Analytic Aggregate Function Support

TDV supports the analytic aggregate functions listed in the table below for Netezza

Netezza Analytic Aggregate Function	Notes
AVG	<p>AVG (NULL) is NULL.</p> <p>Not supported:</p> <ul style="list-style-type: none"> • AVG (BOOLEAN) • AVG (NCHAR) • AVG (NCHAR) DISTINCT • AVG (NVARCHAR) • AVG (NVARCHAR) DISTINCT
COUNT	
MAX	<p>MAX (NULL) is NULL.</p> <p>Not supported:</p> <ul style="list-style-type: none"> • MAX (BOOLEAN)

Netezza Analytic Aggregate Function	Notes
	<ul style="list-style-type: none"> • MAX (BOOLEAN) DISTINCT • MAX (NCHAR) • MAX (NCHAR) DISTINCT • MAX (NVARCHAR) • MAX (NVARCHAR) DISTINCT
MIN	<p>MIN (NULL) is NULL.</p> <p>Not supported:</p> <ul style="list-style-type: none"> • MIN (BOOLEAN) • MIN (BOOLEAN) DISTINCT • MIN (NCHAR) • MIN (NCHAR) DISTINCT • MIN (NVARCHAR) • MIN (NVARCHAR) DISTINCT
SUM	<p>SUM (NULL) is NULL.</p> <p>Not supported:</p> <ul style="list-style-type: none"> • SUM (BOOLEAN) • SUM (BOOLEAN) DISTINCT • SUM (NCHAR) • SUM (NCHAR) DISTINCT • SUM (NVARCHAR) • SUM (NVARCHAR) DISTINCT

Netezza Analytic Function Support

TDV supports the analytic functions listed in the table below for Netezza.

Netezza Analytic Function	Notes
DENSE_RANK	
FIRST_VALUE	
LAG	
LAST_VALUE	
LEAD	
RANK	
ROW_NUMBER	
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
VAR_POP	
VAR_SAMP	
VARIANCE	
VARIANCE_POP	
VARIANCE_SAMP	

Netezza Binary Function Support

TDV supports the binary functions listed in the table below for Netezza .

Netezza Binary Function	Notes
INT1AND, INT2AND, INT4AND, INT8AND	Bitwise AND
INT1NOT, INT2NOT, INT4NOT, INT8NOT	Bitwise NOT
INT1OR, INT2OR, INT4OR, INT8OR	Bitwise OR
INT1SHL, INT2SHL, INT4SHL, INT8SHL	Bitwise shift left
INT1SHR, INT2SHR, INT4SHR, INT8SHR	Bitwise shift right
INT1XOR, INT2XOR, INT4XOR, INT8XOR	Bitwise EXCLUSIVE OR

Netezza Character Function Support

TDV supports the character functions listed in the table below for Netezza.

Netezza Character Function	Notes
ASCII	
BTRIM	
CHR	
CONCAT	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
DLE_DST	
INITCAP	
INSTR	

Netezza Character Function	Notes
LE_DST	
LENGTH	
LOWER	
LPAD	Length limit of 4000.
LTRIM	
POSITION	Not supported: <ul style="list-style-type: none"> • NCHAR argument • NVARCHAR argument
REPEAT	
REPLACE	Not available in Netezza. Netezza has a TRANSLATE function, but it works differently.
RPAD	
RTRIM	
SOUNDEX	
SPACE	Not supported: <ul style="list-style-type: none"> • BIT argument
STRPOS	
SUBSTR	
SUBSTRING	Not supported: <ul style="list-style-type: none"> • NCHAR argument • NVARCHAR argument

Netezza Character Function	Notes
TRANSLATE	
TRIM	
TRUNC	
UNICHR	
UNICODE	
UPPER	

Netezza Conditional Function Support

TDV supports the conditional functions listed in the table below for Netezza.

Netezza Conditional Function	Notes
COALESCE	
DECODE	
NULLIF	Not supported: <ul style="list-style-type: none"> • BIT argument • INTERVAL argument • NCHAR argument • NVARCHAR argument

Netezza Conversion Function Support

TDV supports the conversion functions listed in the table below for Netezza.

Netezza Conversion Function	Notes
CAST	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR first argument • NULL first argument • NVARCHAR first argument • <any_number> AS NULL
TO_CHAR	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
TO_DATE	Not supported: <ul style="list-style-type: none"> • NCHAR argument • NVARCHAR argument
TO_NUMBER	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
TO_TIMESTAMP	Not supported: <ul style="list-style-type: none"> • NCHAR argument • NVARCHAR argument

Netezza Date Function Support

TDV supports the date functions listed in the table below for Netezza.

Netezza Date Function	Notes
CURRENT_DATE	
CURRENT_TIME	
CURRENT_TIMESTAMP	
DAY	
MONTH	
YEAR	

Netezza Numeric Function Support

TDV supports the numeric functions listed in the table below for Netezza.

Netezza Numeric Function	Notes
ABS	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
ACOS	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
ASIN	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument

Netezza Numeric Function	Notes
	NVARCHAR argument
ATAN	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
ATAN2	
CEIL	
CEILING	
COS	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
COT	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
DEGREES	
EXP	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
FACTORIAL	

Netezza Numeric Function	Notes
FLOOR	
LN	
LOG	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
NVL	
NVL2	
PI	
POW	
POWER	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
RADIANS	
RANDOM	
ROUND	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
SIGN	

Netezza Numeric Function	Notes
SIN	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
SQRT	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument
TAN	Not supported: <ul style="list-style-type: none"> • BIT argument • NCHAR argument • NVARCHAR argument

Netezza Phonetic Function Support

TDV supports the phonetic functions listed in the table below for Netezza.

Netezza Phonetic Function	Notes
DBL_MP	
NYSIIS	
PRI_MP	
SCORE_MP	
SEC_MP	

Netezza Statistical Analytic Aggregate Function Support

TDV supports the statistical analytic aggregate functions listed in the table below for Netezza.

Netezza Statistical Analytic Aggregate Function	Notes
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
VAR_POP	
VAR_SAMP	
VARIANCE	
VARIANCE_POP	
VARIANCE_SAMP	

Netezza Time Function Support

TDV supports the time functions listed in the table below for Netezza.

Netezza Time Function	Notes
ADD_MONTHS	
DATE_PART	

Netezza Time Function	Notes
DATE_TRUNC	
EXTRACT	
LAST_DAY	
MONTHS_BETWEEN	
NEXT_DAY	
NOW	
TIMEOFDAY	
TIMESTAMP	

References

Refer to the following Guides for further details about the capabilities of the data source:

Capabilities	Section
Query Engine	User Guide, Chapter <i>TDV Query Engine Optimizations</i>
Data ship	User Guide, Chapter <i>Data Ship Performance Optimization</i>
Caching	User Guide, Chapter <i>TDV Caching</i>
Performance Optimization	User Guide, Chapter <i>Performance Tuning</i>
TDV Massively Parallel Processing Engine	User Guide, Chapter <i>Configuring the TDV MPP Engine</i>

Capabilities	Section
Kerberos	Administration Guide Chapter <i>Configuring Kerberos</i>

TIBCO Product Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join the TIBCO Community.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the [TIBCO Product Documentation](#) website, mainly in HTML and PDF formats.

The [TIBCO Product Documentation](#) website is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

The following documentation for this product is available on the [TIBCO® Data Virtualization](#) page.

- **Users**
 - TDV Getting Started Guide
 - TDV User Guide
 - TDV Web UI User Guide
 - TDV Client Interfaces Guide
 - TDV Tutorial Guide
 - TDV Northbay Example
- **Administration**
 - TDV Installation and Upgrade Guide
 - TDV Administration Guide
 - TDV Active Cluster Guide
 - TDV Security Features Guide
- **Data Sources**

TDV Adapter Guides

TDV Data Source Toolkit Guide (Formerly Extensibility Guide)

- **References**

TDV Reference Guide

TDV Application Programming Interface Guide

- **Other**

TDV Business Directory Guide

TDV Discovery Guide

- *TIBCO TDV and Business Directory Release Notes* Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

How to Contact TIBCO Support

Get an overview of [TIBCO Support](#). You can contact TIBCO Support in the following ways:

- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the [TIBCO Support](#) website.
- For creating a Support case, you must have a valid maintenance or support contract with TIBCO. You also need a user name and password to log in to [TIBCO Support](#) website. If you do not have a user name, you can request one by clicking **Register** on the website.

Release Version Support

TDV 8.5 is designated as a Long Term Support (LTS) version. Some release versions of TIBCO® Data Virtualization products are selected to be long-term support (LTS) versions. Defect corrections will typically be delivered in a new release version and as hotfixes or service packs to one or more LTS versions. See also

https://docs.tibco.com/pub/tdv/general/LTS/tdv_LTS_releases.htm.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote on feature requests from within the [TIBCO Ideas Portal](#). For a free registration, visit [TIBCO Community](#).

Legal and Third-Party Notices

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE “LICENSE” FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of TIBCO Software Inc.

TIBCO, TIBCO logo, TIBCO O logo, ActiveSpaces, Enterprise Messaging Service, Spotfire, TERR, S-PLUS, and S+ are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

This software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the

readme file for the availability of this software version on a specific operating system platform.

THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

This and other products of TIBCO Software Inc. may be covered by registered patents. Please refer to TIBCO's Virtual Patent Marking document (<https://www.tibco.com/patents>) for details.

Copyright © 2002-2023 Cloud Software Group, Inc All Rights Reserved.