



# **TIBCO® Data Virtualization**

## **Teradata Adapter Guide**

Version 8.7.0 | October 2023



# Contents

---

<b>Contents</b>	<b>2</b>
<b>TDV Teradata Adapter</b>	<b>4</b>
Introduction	4
Obtain and Install the Driver for Teradata	4
Datasource Configuration	5
Basic Tab	5
Advanced Tab	9
Teradata Data Source Introspection	26
Teradata Data Source Limitations	27
Teradata Caching Limitations	27
Teradata to TDV Data Types	29
Teradata Cache Mapping	32
Teradata Function Support	34
Teradata Aggregate Function Support	34
Teradata Analytic Function Support	35
Teradata Character Function Support	36
Teradata Conditional Function Support	38
Teradata Conversion Function Support	38
Teradata Date Function Support	39
Teradata Number Function Support	40
Teradata Specific Properties	42
References	44
<b>TIBCO Product Documentation and Support Services</b>	<b>45</b>
How to Access TIBCO Documentation	45
How to Contact TIBCO Support	46
Release Version Support	46



How to Join TIBCO Community .....	47
<b>Legal and Third-Party Notices .....</b>	<b>48</b>



# TDV Teradata Adapter

---

## Introduction

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

[Datasource Configuration](#)

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

[Teradata Data Source Introspection](#)

[Teradata Data Source Limitations](#)

[Data Type Mappings](#)

[Teradata Function Support](#)

[References](#)

## Obtain and Install the Driver for Teradata

TDV is already preconfigured to use the Teradata drivers, but the drivers must be downloaded from Teradata and installed independently.

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

### Obtain and install the JDBC driver for Teradata

Download the Teradata JDBC drivers, which are packaged in a zip or TAR archive. For example, navigate to following Teradata download site:

<https://downloads.teradata.com/download/connectivity/jdbc-driver>

1. Copy the following adapter JARs from the Teradata driver archive:



Teradata	Teradata 14	Teradata 15	Teradata 16
tdgssconfig.jar	tdgssconfig.jar	tdgssconfig.jar	tdgssconfig.jar
terajdbc4.jar	terajdbc4.jar	terajdbc4.jar	terajdbc4.jar

2. Paste the adapter JAR files into this TDV installation directory:  
<TDV\_install\_dir>/conf/adapters/system/teradata\_<version>
3. Restart the TDV Server to initiate use of the new adapter JAR files.

## 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 Teradata is 1025.
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.
<a href="#">Pass-through Login</a>	Flag to indicate whether pass-through login is enabled or not.
<a href="#">Transaction Isolation</a>	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 Uncommitted—Dirty reads, nonrepeatable reads, and phantom reads can occur.
- Read Committed—Nonrepeatable reads and phantom reads can occur.
- Repeatable Read—Only 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.

## 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.
<a href="#">Concurrent Request Limit</a>	Works with the Massively Parallel Processing engine configuration parameters to control the amount of parallelization for the queries for a particular data source.
<a href="#">Use X Views</a>	Returns data only for rows containing information on objects that the requesting user owns, created, has privileges on, or has been granted access through a current or nested role.
<a href="#">Query Banding</a>	<p>Turns the query banding feature on (At Session Level) or off (Off). Stores query context information in the Teradata session table so it can be recovered after a system reset. Query banding takes effect when the data source is next used.</p> <p>Connection pooling has no effect on query band data.</p>
<a href="#">QueryBand Properties</a>	<p>Click QueryBand Properties on the right to open a dialog box in which to specify property-value pairs to store in the session table.</p> <p>Four properties are available by default. For the first three, if the default value (including brackets) appears alone in the value field, it is replaced with the actual value at run time.</p> <ul style="list-style-type: none"> <li>• TDV_USER. ID of the TDV user, application or report that originated the request from a middle-tiered application. The default value is &lt;TDV_USER&gt;.</li> <li>• DOMAIN. The default value is &lt;DOMAIN&gt;.</li> <li>• SESSION_NAME. The default value is &lt;SESSION_NAME&gt;.</li> <li>• SYS. The default value is TDV.</li> </ul> <p>TDV administrator can click Add property to add custom</p>



	name-value pairs.
<a href="#">Is dataship source</a>	Indicates whether the physical data source might be used as a source of shipped tables to another data ship enabled data source
<a href="#">Is dataship target</a>	Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.
<a href="#">Lower bound/Upper bound for dataship</a>	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.
<a href="#">Schema Path for Temp Tables</a>	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.
<a href="#">Temp Table Prefix</a>	A character string addition to temporary table names so that they are recognized if they are needed.
<a href="#">Lower bound/Upper bound for dataship</a>	Use the FastLoad or FastExport utility to speed up queries. Cardinality information determines whether to use Fastpath or JDBC default loading for a given query.
<a href="#">FastExport Session Count</a>	The number of FastExport sessions to use.
<a href="#">FastLoad Session Count</a>	The number of FastLoad sessions to use.

## 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

## Use X Views

Returns data only for rows containing information on objects that the requesting user owns, created, has privileges on, or has been granted access through a current or nested role.

## Data type

Bool

## Default Value

False

## Remarks

None

## Query Banding

Turns the query banding feature on (At Session Level) or off (Off). Stores query context information in the Teradata session table so it can be recovered after a system reset. Query banding takes effect when the data source is next used.

Connection pooling has no effect on query band data.

## Data Type

String



## Default Value

Off

## Remarks

None

## QueryBand Properties

Click QueryBand Properties on the right to open a dialog box in which to specify property-value pairs to store in the session table.

Four properties are available by default. For the first three, if the default value (including brackets) appears alone in the value field, it is replaced with the actual value at run time.

- TDV\_USER. ID of the TDV user, application or report that originated the request from a middle-tiered application. The default value is <TDV\_USER>.
- DOMAIN. The default value is <DOMAIN>.
- SESSION\_NAME. The default value is <SESSION\_NAME>.
- SYS. The default value is TDV.

TDV administrator can click Add property to add custom name-value pairs.

## Data Type

String

## Default Value

""

## Remarks

None



## Is dataship source

Indicates whether the physical data source might be used as a source of shipped tables to another data ship enabled data source.

### Data Type

Bool

### Default Value

False

### 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.

### Data Type

Bool

### Default Value

False

### Remarks

To make changes in this field, Is dataship source must be enabled.

## 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

To make changes in this field, Is dataship source must be enabled.

## 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

To make changes in this field, Is dataship target must be enabled.

## 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

To make changes in this field, Is dataship target must be enabled.

## Enable FastLoad/FastExport for large tables

Use the FastLoad or FastExport utility to speed up queries. Cardinality information determines whether to use Fastpath or JDBC default loading for a given query.

## Data Type

Bool

## Default Value

False

## Remarks

None

## FastExport Session Count

The number of FastExport sessions to use.

## Data Type

Numeric



## Default Value

4

## Remarks

To make changes in this field, Enable FastLoad/FastExport for large tables field must be enabled.

## FastLoad Session Count

The number of FastLoad sessions to use.

## Data Type

Numeric

## Default Value

4

## Remarks

To make changes in this field, Enable FastLoad/FastExport for large tables field must be enabled.

## 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.

# Teradata Data Source Introspection

Introspection of Teradata data sources has these characteristics:



- Introspection of a Teradata RDBMS data sources reveals tables, views, procedures, functions, and macros.
- Introspected Teradata functions and macros are displayed as procedures in the Studio resource tree.
- Teradata macro parameters map to TDV parameters, with the appropriate TDV data types and references for inputs and results.

## Teradata Data Source Limitations

The following data types cannot be used as part of an INDEX:

- CLOB
- BLOB
- XML

## Teradata Caching Limitations

Teradata has a known issue that affects data caching for TDV. Because of this issue, you might be unable to cache data, and you might get incorrect query results against Teradata when Ignore Trailing Spaces is set to FALSE in TDV. The issue is caused by the Teradata driver's management of character data when using UTF-8 character sets.

To solve both the caching and the query problems, you can do one of the following:

- Change the global server setting for Ignore Trailing Spaces to true. (In Studio, choose Administration > Configuration. Locate and select Ignore Trailing Spaces, and click True for Value.)
- Change the Teradata connection string to use UTF-16 by substituting CHARSET=UTF16 for CHARSET=UTF8, and save the data source. Then recreate the cache\_status table and refresh the cache.
- Change the Teradata connection string to use ASCII by substituting CHARSET=ASCII for CHARSET=UTF8, and save the data source. Then recreate the cache\_status table and refresh the cache. This solution does not work for data that contains multi-byte international characters because the characters are not saved or retrieved correctly.



To solve just the caching problem, cache data in a different data source (instead of Teradata).

To solve just the query problem, provide query hints (see the Section *Specifying Query Hints* in the User Guide) on queries against Teradata where filters are on CHAR columns:

```
{ OPTION IGNORE_TRAILING_SPACES="True" }
```

## Teradata Multi-Table Caching Limitations

The Teradata Fast Export and Fast Load features are supported for caching with the TDV multi-table caching option. The cache must have no duplicate rows of data and be configured as specified in See the Section *Configuring Teradata for Use as a Multi-table Cache Target* in the User Guide

- Teradata FastLoad requires that the target table be empty.
- Teradata has limitation on how many concurrent FastLoad and FastExport tasks can run in parallel. Parallelism is controlled by the MaxLoadTasks and MaxLoadAWT parameters. Any FastLoad task exceeding the limitation is rejected.
- For Teradata, the maximum session for each FastLoad or FastExport job is limited to the number of AMPs of the Teradata database. Typically, eight sessions work well for most scenarios.
- For Teradata, a row fetch size bigger than 64 KB causes a Teradata error. Teradata big objects can be configured using Teradata to return data in differed transfer mode. Refer to your Teradata documentation to determine the best solution for you if you have data rows that return 64 KB or greater of data.
- The following data type and function support restrictions exist.

Data Source	Cache Target	Data Types Not Supported	Functions Not Supported
Oracle	Teradata	BLOB, CLOB, LONG, LONGRAW, NCLOB	No results returned after refreshing the cache against INTERVALDAYTOSECOND and INTERVALYEARTOMONTH.
SQL Server 2008	Teradata	BINARY, IMAGE, NTEXT,	



Data Source	Cache Target	Data Types Not Supported	Functions Not Supported
		TEXT, VARBINARY	
Sybase	Teradata	BINARY, IMAGE, TEXT, VARBINARY	
Teradata	Teradata	BYTE, BLOB, CLOB, LONGVARCHAR	
Vertica 5.0 and 6.1	Teradata	BINARY, VARBINARY	

## TDV Native Loading Option Teradata Limitation

If the TDV native load option is active and the data identified to be moved to the cache has one or more duplicate rows, Teradata 13 will allow the duplicate rows.

## Teradata to TDV Data Types

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

Teradata data types have these characteristics:

- FLOAT and REAL data types are synonymous with DOUBLE PRECISION.
- For Teradata version 15, the maximum length for BINARY and VARBINARY is 64000; for CHAR and VARCHAR it is 32000; for BLOB it is 2097088000; for CLOB it is 1048544000; and for JSON it is 8388096.
- For all supported versions of Teradata except version 15, the maximum length for BINARY, CHAR, VARBINARY, and VARCHAR is 32000; and the JSON data type is not supported.
- IN operator with subquery cannot be pushed down,
- The native data types BLOB, CLOB, JSON, and XML are not supported in DISTINCT, EXCEPT, GROUP BY, HAVING, INTERSECT, JOIN ON, ORDER BY or UNION clauses,



Teradata Data Type	TDV Data Type
BIGINT	BIGINT
BLOB	BLOB
BYTE	BINARY
BYTEINT	TINYINT
CHAR	CHAR
CLOB	CLOB
DATE	DATE
DECIMAL	DECIMAL
DOUBLE PRECISION	DOUBLE
FLOAT	DOUBLE
GRAPHIC	CHAR
INTEGER	INTEGER
INTERVAL DAY	INTERVAL DAY
INTERVAL DAY TO HOUR	INTERVAL DAY TO HOUR
INTERVAL DAY TO MINUTE	INTERVAL DAY TO MINUTE
INTERVAL DAY TO SECOND	INTERVAL DAY TO SECOND
INTERVAL HOUR	INTERVAL HOUR
INTERVAL HOUR TO MINUTE	INTERVAL HOUR TO MINUTE
INTERVAL HOUR TO SECOND	INTERVAL HOUR TO SECOND



Teradata Data Type	TDV Data Type
INTERVAL MINUTE	INTERVAL MINUTE
INTERVAL MINUTE TO SECOND	INTERVAL MINUTE TO SECOND
INTERVAL MONTH	INTERVAL MONTH
INTERVAL SECOND	INTERVAL SECOND
INTERVAL YEAR	INTERVAL YEAR
INTERVAL YEAR TO MONTH	INTERVAL YEAR TO MONTH
LONG VARCHAR	CLOB
NUMERIC	Arbitrary NUMERIC
PERIOD(DATE)	VARCHAR
PERIOD(TIME)	VARCHAR
PERIOD(TIMESTAMP)	VARCHAR
PERIOD(TIME WITH TIME ZONE)	VARCHAR
PERIOD(TIMESTAMP WITH TIME ZONE)	VARCHAR
REAL	DOUBLE
SMALLINT	SMALLINT
TIME	TIME
TIME WITH ZONE	VARCHAR
TIMESTAMP	TIMESTAMP



Teradata Data Type	TDV Data Type
TIMESTAMP WITH ZONE	VARCHAR
VARBYTE	VARBINARY
VARCHAR	VARCHAR
VARGRAPHIC	VARCHAR
XML	XML (version 15 and above)

## Teradata Cache Mapping

This section discusses the data type mappings and restrictions for caches stored on Teradata. Data type overrides for version 13 is indicated in square brackets. Data types not listed in the table cannot be cached.

Data Type	Preferred Data Type	Other Allowed Native Types
BIGINT	BIGINT [12, 13] CHAR(20)	DECIMAL(19+,0), VARCHAR(20+)
BINARY(n)	BYTE(n), BLOB [if n > 32,000]	BYTE(n+)
BIT	BYTEINT	DECIMAL(1+,0), larger INTEGER types
BLOB	BLOB	
BOOLEAN	BYTEINT	SMALLINT, INTEGER
CHAR(n)	CHAR(n); CLOB [if n > 32,000]	CHAR(n+), GRAPHIC(n+), VARCHAR(n+), CLOB
CLOB	CLOB	Teradata 15 JDBC driver does not support CLOB column with



Data Type	Preferred Data Type	Other Allowed Native Types
		NULL values when caching to Teradata 15.
DATE	DATE	VARCHAR(10+)
DECIMAL(p,s)	DECIMAL(p,s); CLOB [if p > 18]	DECIMAL(p+,s+), VARCHAR (p+3+), GRAPHIC(p+3+), CLOB
DOUBLE	FLOAT	VARCHAR(24+)
FLOAT	FLOAT	VARCHAR(24+)
INTEGER	INTEGER	DECIMAL(10+,0), VARCHAR(20+), GRAPHIC(20+), , VARGRAPHIC (20+)
NUMERIC(p,s)	DECIMAL(p,s); CLOB [if p > 18]	DECIMAL(p+,s+), VARCHAR (p+3+), GRAPHIC(p+3+), CLOB
OTHER	[cannot be cached]	
SMALLINT	SMALLINT	DECIMAL(5+,0), larger INTEGER types, VARCHAR(20+), VARGRAPHIC(20+)
TIME	VARCHAR(15)	VARCHAR(15+)
TIMESTAMP	TIMESTAMP	VARCHAR(26+)
TINYINT	BYTEINT	DECIMAL(3+,0), larger INTEGER types, VARCHAR(20+), VARGRAPHIC(20+)
VARBINARY(n)	VARBYTE(n); BLOB [if n > 32,000]	VARBYTE(n+)
VARCHAR(n)	VARCHAR(n);	VARCHAR(n+), VARGRAPHIC(n+)



Data Type	Preferred Data Type	Other Allowed Native Types
	CLOB [if n > 32,000]	
XML	CLOB	VARCHAR(*) [Truncates data if column too small]

## Teradata Function Support

TDV supports the following types of functions for Teradata:

- [Teradata Aggregate Function Support](#)
- [Teradata Analytic Function Support](#)
- [Teradata Character Function Support](#)
- [Teradata Conditional Function Support](#)
- [Teradata Conversion Function Support](#)
- [Teradata Date Function Support](#)
- [Teradata Number Function Support](#)

## Teradata Aggregate Function Support

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

Teradata Aggregate Function	Notes
AVG	
COUNT	Large objects (BLOB, CLOB, and so on) not supported.
MAX	Decimal not supported (may introduce rounding error).



Teradata Aggregate Function	Notes
MEDIAN	Version 15 and above.
MIN	Decimal not supported (may introduce rounding error).
SUM	

## Teradata Analytic Function Support

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

Teradata Analytic Function	Notes
AVG	
COUNT	Large objects (BLOB, CLOB, and so on) not supported. Version 15: JSON not supported.
CUME_DIST	Version 15 and above.
DENSE_RANK	Version 15 and above.
FIRST_VALUE	Includes ignore nulls. Version 15 and above.
LAST_VALUE	Includes ignore nulls. Version 15 and above.
MAX	Decimal not supported (may introduce rounding error). Version 15: JSON not supported.
MIN	Decimal not supported (may introduce rounding error). MIN (date) incorrectly returns NULL if one of the date columns is NULL. Version 15: JSON not supported.
PERCENT_RANK	



Teradata Analytic Function	Notes
PERCENTILE_CONT	Version 15 and above.
PERCENTILE_DIST	Version 15 and above.
RANK	
ROW_NUMBER	
STDDEV_POP	
STDDEV_SAMP	
SUM	
VAR_POP	
VAR_SAMP	
VARIANCE_POP	
VARIANCE_SAMP	
LAG	Version 16
LEAD	Version 16

## Teradata Character Function Support

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

Teradata Character Function	Notes
ASCII	Version 15 and above.
CHAR_LENGTH	



Teradata Character Function	Notes
CHR	Version 15 and above.
CONCAT	
INITCAP	Version 15 and above.
INSTR	Version 15 and above.
LEAD	Version 15 and above.
LEFT	Version 15 and above.
LENGTH	
LOWER	
LPAD	Version 15 and above.
LTRIM	Version 15 and above.
POSITION	
REPLACE	Version 15 and above.
REVERSE	Version 15 and above.
RIGHT	Version 15 and above.
RPAD	Version 15 and above.
RTRIM	
SPACE	Not supported.
SUBSTRING	
TRANSLATE	Version 15 and above.



Teradata Character Function	Notes
TRIM	
UPPER	

## Teradata Conditional Function Support

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

Teradata Conditional Function	Notes
COALESCE	
DECODE	Version 15 and above.
GREATEST	Version 15 and above.  From Version 16, the DATE and TIMESTAMP data types are accepted as input and output.
LEAST	Version 15 and above.  From Version 16, the DATE and TIMESTAMP data types are accepted as input and output.
NULLIF	
NVL	
NVL2	Version 15 and above.

## Teradata Conversion Function Support

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



Teradata Conversion Function	Notes
CAST	
FORMAT	
FORMAT_DATE	
PARSE_DATE	
PARSE_TIME	
PARSE_TIMESTAMP	
TO_CHAR	Version 15 and above. Output format may differ for push versus no-push, but precision and scale are the same. For example, Teradata 15 TDV might return 1.234567890000000E 005, while TDV (DISABLE_PUSH='TRUE') returns 123456.789.
TO_DATE	
TO_TIMESTAMP	Version 15 and above.
TO_TIMESTAMP_TZ	Version 15 and above.

## Teradata Date Function Support

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

Teradata Date Function	Notes
CURRENT_DATE	
CURRENT_TIME	
CURRENT_TIMESTAMP	



Teradata Date Function	Notes
DAY	
EXTRACTDAY	
EXTRACTHOUR	Not supported: EXTRACTHOUR from INTERVAL_DAY or INTERVAL_MINUTE.
EXTRACTMINUTE	
EXTRACTMONTH	
EXTRACTSECOND	
EXTRACTYEAR	
LAST_DAY	Version 15 and above.
MONTH	
MONTHS_BETWEEN	Version 15 and above.
NEXT_DAY	Version 15 and above.
YEAR	

## Teradata Number Function Support

TDV supports the number functions listed in the table below for Teradata, except as marked.

Teradata Number Function	Notes
ABS	
ACOS	



Teradata Number Function	Notes
ASIN	
ATAN	
CEILING	Versions 13 and above.
COS	
COT	
DEGREE	Not supported.
EXP	
FLOOR	Versions 13 and above.
LOG	
PI	Not supported.
RADIANS	Not supported.
ROUND	Version 15 and above.  From Version 16, TIMESTAMP data type is returned as output if the same type is used as input.
SIGN	Version 15 and above.
SIN	
SQRT	
TAN	
TRUNC	Version 15 and above.



## Teradata Specific Properties

This section describes the connection properties that are specific to the Teradata data source.

Port	<p>Port number for the data source to connect with the host.</p> <p>The Port number for Teradata is 1025.</p>
Connection URL Pattern	<p>Sends debug log messages to system.out rather than to a file. The pattern is jdbc:teradata://&lt;HOST&gt;/DBS_PORT=&lt;PORT&gt;/DATABASE=&lt;DATABASE_NAME&gt;/CHARSET=UTF8,COMPAT_DBS=true. To view debug messages, append ,LOG=DEBUG (including the initial comma) to the pattern.</p>
Enable FastLoad/FastExport for large tables	<p>Use the FastLoad or FastExport utility to speed up queries. Cardinality information determines whether to use Fastpath or JDBC default loading for a given query.</p>
FastExport Session Count	<p>The number of FastExport sessions to use.</p>
FastLoad Session Count	<p>The number of FastLoad sessions to use.</p>
Query Banding	<p>Turns the query banding feature on (At Session Level) or off (Off). Stores query context information in the Teradata session table so it can be recovered after a system reset. Query banding takes effect when the data source is next used.</p> <p>Connection pooling has no effect on query band data.</p>
QueryBand Properties	<p>Click QueryBand Properties on the right to open a dialog box in which to specify property-value pairs to store in the session table.</p> <p>Four properties are available by default. For the first three, if the default value (including brackets) appears</p>



	<p>alone in the value field, it is replaced with the actual value at run time.</p> <ul style="list-style-type: none"> <li>• TDV_USER. ID of the TDV user, application or report that originated the request from a middle-tiered application. The default value is &lt;TDV_USER&gt;.</li> <li>• DOMAIN. The default value is &lt;DOMAIN&gt;.</li> <li>• SESSION_NAME. The default value is &lt;SESSION_NAME&gt;.</li> <li>• SYS. The default value is TDV.</li> </ul> <p>TDV administrator can click Add property to add custom name-value pairs.</p>
Use X Views	Returns data only for rows containing information on objects that the requesting user owns, created, has privileges on, or has been granted access through a current or nested role.
Is dataship source	Indicates whether the physical data source might be used as a source of shipped tables to another data ship enabled data source
Lower bound/Upper bound	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.
Is dataship target	Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.
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.
-------------------	--

---

## References

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

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



# 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](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.