



TIBCO® Data Virtualization

SQL Server Adapter Guide

Version 8.7.0 | October 2023

Contents

Contents	2
TDV SQL Server Adapter	4
Introduction	4
Datasource Configuration	4
Basic Tab	4
Advanced Tab	10
Obtain and Install the Driver for SQL Server	27
Connecting to SQL Server datasource using an Active directory Login	28
Microsoft SQL Server 2008 Limitation	28
Data Type Mappings	29
Microsoft SQL Server to TDV Data Types	29
Microsoft SQL Server Cache Mapping	31
Microsoft SQL Server Function Support	33
Microsoft SQL Server Aggregate Function Support	34
Microsoft SQL Server Analytic Aggregate Function Support	34
Microsoft SQL Server Analytic Function Support	35
Microsoft SQL Server Character Function Support	36
Microsoft SQL Server Conditional Function Support	37
Microsoft SQL Server Conversion Function Support	38
Microsoft SQL Server Date Function Support	39
Microsoft SQL Server Encryption Function Support	40
Microsoft SQL Server Numeric Function Support	40
Microsoft SQL Server Time Function Support	42
References	42
TIBCO Product Documentation and Support Services	43
How to Access TIBCO Documentation	43

How to Contact TIBCO Support	44
Release Version Support	44
How to Join TIBCO Community	45
Legal and Third-Party Notices	46

TDV SQL Server Adapter

Introduction

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

[Datasource Configuration](#)

[Obtain and Install the Driver for SQL Server](#)

[Connecting to SQL Server datasource using an Active directory Login](#)

[Microsoft SQL Server 2008 Limitation](#)

[Data Type Mappings](#)

[Microsoft SQL Server Function Support](#)

[References](#)

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.
Named Instance	<p>A named instance is identified by the network name of the computer and the instance name that you specify during installation. You must specify Named instance or port when connecting. This field is available in Microsoft SQL Server 2019 datasource and above.</p> <p>This is an optional field.</p>
Port	The Port number.
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.
Authentication Type	The type of Authentication used by the datasource.
Keytab File	Use this option to enable Kerberos security through Keytab files. Type the full path to the Keytab file.
Principal Name	Indicates the Kerberos Service Principal Name.

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

""

Named Instance

A named instance is identified by the network name of the computer and the instance name that you specify during installation. You must specify Named instance or port when connecting. This is an optional field. This field is available in Microsoft SQL Server 2019 datasource and above.

Data type

String

Default Value

""

Remarks

This field is available in Microsoft SQL Server 2019 datasource.

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.

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.

Keytab File

Use this option to enable Kerberos security through Keytab files. Type the full path to the Keytab file.

Data type

String

Default Value

""

Remarks

This field is available only if you choose KERBOS Authentication type.

Principal Name

Indicates the Kerberos Service Principal Name.

Data Type

String

Default Value

""

Remarks

This field is available only if you choose KERBOS Authentication type.

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 Minimum Size	Minimum number of connections in the pool even when the pool is inactive.
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 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.
Select Mode	<p>Direct or Cursor</p> <p>Direct—Sends all results to the adapter in one request. Each statement establishes its own connection to the database using the same connection properties as the original connection, with auto-commit enabled. Java Transaction API (JTA) is not supported. Direct mode does not support operations where the adapter creates a second statement internally. A typical exception message is “Cannot start a cloned connection while in manual transaction mode.”</p> <p>Cursor—Allows you to work with a smaller set of rows that are returned by the SQL statements.</p>

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.
Show All Database	<p>Check to list all databases accessible using these credentials during introspection.</p> <p>If a SQL Server database is off-line for the instance you</p>

	are attempting to introspect, you might see NPE exceptions in the log file.
Ignore Procedure Return Parameter	Check this option to ignore the procedure return parameter.
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.
Is dataship target	Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.
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.
Enable Bulk Import/Export	Related to the data ship feature capability. Several fields are available only if others are checked. For details, see the section Data Ship Performance Optimizations in the User Guide.

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

Show All Database

Check to list all databases accessible using these credentials during introspection.

If a SQL Server database is off-line for the instance you are attempting to introspect, you might see NPE exceptions in the log file.

Data Type

Bool

Default Value

False

Remarks

None

Ignore Procedure Return Parameter

Check this option to ignore the procedure return parameter.

Data Type

Bool

Default Value

False

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 Bulk Import/Export

Related to the data ship feature capability.

Several fields are available only if others are checked. For details, see the section Data Ship Performance Optimizations in the User Guide.

Data Type

Bool

Default Value

False

Remarks

To make changes in this field, Is dataship source 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

Select Mode

Direct or Cursor

Direct—Sends all results to the adapter in one request. Each statement establishes its own connection to the database using the same connection properties as the original connection, with auto-commit enabled. Java Transaction API (JTA) is not supported. Direct mode does not support operations where the adapter creates a second statement internally. A typical exception message is “Cannot start a cloned connection while in manual transaction mode.”

Cursor—Allows you to work with a smaller set of rows that are returned by the SQL statements.

Data Type

String

Default Value

Direct

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.

Obtain and Install the Driver for SQL Server

Obtain and install the JDBC driver for SQL Server

Visit the following URL for the Microsoft JDBC Driver for SQL Server Web page and download the JDBC driver for your version of SQL Server:

<http://msdn.microsoft.com/en-us/data/aa937724.aspx>

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

1. Follow the instructions provided to extract the files for your version of SQL Server.
2. Locate the JAR file for your version.

Version	JAR File
2008, 2012, 2014, 2016	sqljdbc4.jar or sqljdbc_<version>_<language>.tar.gz
2019	mssql-jdbc-7.4.1.jre11.jar mssql-jdbc-7.4.1.jre11.jar

3. Put the JAR file in the directory appropriate to the version:

```
<TDV_install_dir>\conf\adapters\system\microsoft_sql_server_<ver>
```

For UNIX, navigate to the directory where you want the driver unpacked, and type:

```
gzip -d sqljdbc_<version>_<language>.tar.gz.
```

4. Restart the TDV Server.
5. If integrated security is required, it is recommended the you keep the SQL server JDBC drivers in a common location. For example, <TDV_install_dir>\apps\dml\cis_ds_mssql\lib.

Connecting to SQL Server datasource using an Active directory Login

Follow these steps to connect to a SQL Server datasource using an Active Directory login:

1. Download the Microsoft jTDS JDBC driver
2. From Studio, create a new data source
3. Select "New Adapter" using the "Microsoft SQL Server <desired version>" parent adapter
4. Modify the "Adapter Class Name" to: `net.sourceforge.jtds.jdbc.Driver`
5. Modify the Connection URL Pattern to:
`jdbc:jtds:sqlserver://<HOST>:<PORT>/<DATABASE_NAME>;tds=8.0;lastupdatecount=true`
6. Copy the jTDS driver to `/conf/adapters/custom/<name of the new adapter>`
7. Restart TDV
8. Create a new data source from the new adapter
9. Supply the following details in the data source panel:
 - Host:
 - Port:
 - Database Name:
 - Login: <user name (no domain)>
 - Password: <user's password>
10. In the Advanced tab, add a JDBC Connection Property:
 - Property: domain
 - Value: <domain name of AD user>

Microsoft SQL Server 2008 Limitation

There is a known problem with Microsoft related to JRE that results in disabling of the 3DES_EDE_CBC transport layer security algorithm. If you encounter this problem, you can enable 3DES_EDE_CBC in `<TDV_install_dir>/jdk/conf/security/java.security`.

To re-enable 3DES_EDE_CBC

Navigate to the <TDV_install_dir>/jdk/conf/security/java.security.

1. Open the file and remove 3DES_EDE_CBC from the jdk.tls.disabledAlgorithms setting.
2. Restart the TDV Server.

Data Type Mappings

Microsoft SQL Server to TDV Data Types

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

Microsoft SQL Server Data Type	TDV Data Type	Notes
BIGINT	BIGINT	
BINARY	BINARY	
BIT	BIT	
CHAR	CHAR	
DATE	DATE	SQL Server 2008 and 2012.
DATETIME	TIMESTAMP	
DATETIME2	TIMESTAMP	SQL Server 2008 and 2012.
DATETIME2(0) – DATETIME2(7)	TIMESTAMP	SQL Server 2008 and 2012.

Microsoft SQL Server Data Type	TDV Data Type	Notes
DATETIMEOFFSET	VARCHAR	SQL Server 2008 and 2012.
DATETIMEOFFSET(0) – DATETIMEOFFSET(7)	VARCHAR	SQL Server 2008 and 2012.
DECIMAL	DECIMAL	
FLOAT	DOUBLE	
IMAGE	BLOB	
INT	INTEGER	
INT IDENTITY	INTEGER	
MONEY	DECIMAL	
NCHAR	CHAR	
NTEXT	CLOB	
NUMERIC	NUMERIC	
NVARCHAR	VARCHAR	
REAL	FLOAT	
SMALLDATETIME	TIMESTAMP	
SMALLINT	SMALLINT	
SMALLMONEY	DECIMAL	
SQL_VARIANT	OTHER	ODBC does not fully support this data type.

Microsoft SQL Server Data Type	TDV Data Type	Notes
TABLE	OTHER	
TEXT	CLOB	
TIME	TIME	SQL Server 2008 and 2012.
TIME() – TIME(7)	TIME	SQL Server 2008 and 2012.
TIMESTAMP	VARBINARY	
TINYINT	SMALLINT	
UNIQUEIDENTIFIER	CHAR	
VARBINARY	VARBINARY	
VARCHAR	VARCHAR	
XML	XML	

Microsoft SQL Server Cache Mapping

This section discusses the data type mappings and restrictions for caches stored on Microsoft SQL Server 2000, 2005, or 2008. Overrides for the 2008 version are indicated in square brackets.

- SQL Server's page size limits the number of bytes that can be stored directly in a column—so executing DDL causes an error if the resulting table requires a row size greater than this limit. The solution is to either raise the page size for the database, or to use indirect storage types such as TEXT and IMAGE. TDV chooses TEXT and IMAGE types if a value requires more than 255 bytes of storage for this reason, although SQL Server does allow VARCHAR and VARBINARY up to 8,000 bytes. Hand-tuning of the data types used in a table can improve storage efficiency.
- Microsoft SQL Server TINYINT has a range 0 to 255, and TDV TINYINT is -128 to 127, so these types are not compatible.

- DATETIME has only 3.33ms accuracy, so rounding error may occur.

Data Type	Preferred Native Type	Other Allowed Native Types
BIGINT	BIGINT	DECIMAL(19+,0), larger INTEGER types, VARCHAR(20+), NVARCHAR(20+)
BINARY(n)	BINARY(n); IMAGE [if n > 255]	BINARY(n+), IMAGE
BIT	BIT	DECIMAL(1+,0), larger INTEGER types
BLOB	IMAGE	
BOOLEAN	BIT	TINYINT, SMALLINT, INTEGER, BIGINT
CHAR(n)	CHAR(n); TEXT [if p > 38]	CHAR(n+), NCHAR(n+), VARCHAR(n+), NVARCHAR(n+), TEXT, NTEXT
CLOB	TEXT	NTEXT
DATE	DATE [2008] VARCHAR(10)	VARCHAR(10+)
DECIMAL(p,s)	DECIMAL(p,s); TEXT [if p > 38]	DECIMAL(p+,s+), VARCHAR (p+3+), NVARCHAR(n+), TEXT, NTEXT
DOUBLE	FLOAT	VARCHAR(24+)
FLOAT	REAL	FLOAT, VARCHAR(24+)
INTEGER	INTEGER	DECIMAL(10+,0), larger INTEGER types, VARCHAR(20+), NVARCHAR(20+)

Data Type	Preferred Native Type	Other Allowed Native Types
NUMERIC(p,s)	DECIMAL(p,s); TEXT [if p > 38]	DECIMAL(p+,s+), VARCHAR (p+3+), NVARCHAR(p+3+), TEXT, NTEXT
OTHER	[cannot be cached]	
SMALLINT	SMALLINT	DECIMAL(5+,0), larger INTEGER types, VARCHAR(20+), NVARCHAR(20+)
TIME	TIME [2008] VARCHAR(15)	VARCHAR(15+)
TIMESTAMP	DATETIME2 [2008] DATETIME	
TINYINT	SMALLINT	DECIMAL(3+,0), larger INTEGER types, VARCHAR(20+), NVARCHAR(20+)
VARBINARY(n)	VARBINARY(n); IMAGE [if n > 255]	VARBINARY(n+), IMAGE
VARCHAR(n)	VARCHAR(n); TEXT [if n > 255]	VARCHAR(n+), NVARCHAR(n+), TEXT, NTEXT
XML	TEXT	VARCHAR(*) [Truncates data if column is too small], TEXT

Microsoft SQL Server Function Support

TDV supports the following types of functions for Microsoft SQL Server:

- [Microsoft SQL Server Aggregate Function Support](#)
- [Microsoft SQL Server Analytic Function Support](#)

- [Microsoft SQL Server Analytic Aggregate Function Support](#)
- [Microsoft SQL Server Character Function Support](#)
- [Microsoft SQL Server Conditional Function Support](#)
- [Microsoft SQL Server Conversion Function Support](#)
- [Microsoft SQL Server Date Function Support](#)
- [Microsoft SQL Server Encryption Function Support](#)
- [Microsoft SQL Server Numeric Function Support](#)
- [Microsoft SQL Server Time Function Support](#)

Microsoft SQL Server Aggregate Function Support

TDV supports the aggregate functions listed in the table below for Microsoft SQL Server.

Microsoft SQL Server Aggregate Function	Notes
AVG	Unique identifiers not supported.
COUNT	Unique identifiers not supported.
MAX	
MIN	
SUM	Unique identifiers not supported.

Microsoft SQL Server Analytic Aggregate Function Support

TDV supports the analytic aggregate functions listed in the table below for Microsoft SQL Server.

Microsoft SQL Server Analytic Aggregate Function	Notes
CUM_DIST	SQL Server 2012 and 2014.
FIRST_VALUE	SQL Server 2012 and 2014.
LAG	SQL Server 2012 and 2014.
LAST_VALUE	SQL Server 2012 and 2014.
LEAD	SQL Server 2012 and 2014.
PERCENTILE_CONT	SQL Server 2012 and 2014.
PERCENTILE_DISC	SQL Server 2012 and 2014.
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
VARIANCE	
VARIANCE_POP	
VARIANCE_SAMP	

Microsoft SQL Server Analytic Function Support

TDV supports the analytic functions listed in the table below for Microsoft SQL Server 2005 and 2008.

Microsoft SQL Server Analytic Function	Notes
AVG	
COUNT	
DENSE_RANK	
MAX	
MIN	
NTILE	
RANDOM	
RANK	
ROW_NUMBER	
STDDEV	
SUM	
VAR_POP	
VARIANCE	

Microsoft SQL Server Character Function Support

TDV supports the character functions listed in the table below for Microsoft SQL Server.

Microsoft SQL Server Character Function	Notes
ASCII	•

Microsoft SQL Server Character Function	Notes
CONCAT	<ul style="list-style-type: none"> Unique identifiers not supported. When the input timestamp value has no fractional seconds (hh:mm:ss), TDV does not print the fractional part. This is the way many data sources handle this situation. However, SQL Server <i>does</i> add the fractional part (hh:mm:ss.fff).
LENGTH	
LOWER	
POSITION	
REPLACE	
RTRIM	
SOUNDEX	
SPACE	
SUBSTRING	
TRIM	
UPPER	
CHARINDEX	

Microsoft SQL Server Conditional Function Support

TDV supports the conditional functions listed in the table below for Microsoft SQL Server.

Microsoft SQL Server Conditional Function	Notes
COALESCE	
DECODE	Mapped to CASE.
ISNULL	
ISNUMERIC	
NULLIF	NULL literal cannot be the first argument to NULLIF function. NULLIF does not support IMAGE, NTEXT, or TEXT.
NVL	

Microsoft SQL Server Conversion Function Support

TDV supports the conversion functions listed in the table below for Microsoft SQL Server. These conversion functions do not support unique identifiers.

Microsoft SQL Server Conversion Function	Notes
CAST	<ul style="list-style-type: none"> TINYINT is cast to SMALLINT. Cannot cast a number to DOUBLE. Cannot cast a string to any integer data type. Cannot cast floating point to DOUBLE.
FORMAT_DATE	
PARSE_TIMESTAMP	
TO_CHAR	When using this function,TDV pushes down the CONVERT function to SQL Server.

Microsoft SQL Server Conversion Function	Notes
TO_DATE	
TO_NUMBER	
TO_TIMESTAMP	

Microsoft SQL Server Date Function Support

TDV supports the date functions listed in the table below for Microsoft SQL Server.

Microsoft SQL Server Date Function	Notes
CURRENT_DATE	
CURRENT_TIMESTAMP	
DATEADD	SQL Server 2014 only.
DATEDIFF	SQL Server 2008, 2012, 2014. DATEDIFF (DATEPART, STARTDATE, ENDDATE)
DATEPART	SQL Server 2014 only.
DAY	
DAYNAME	SQL Server 2014 only.
DAYOFMONTH	SQL Server 2014 only.
DAYOFWEEK	SQL Server 2014 only.
HOURL	SQL Server 2014 only.
MINUTE	SQL Server 2014 only.

Microsoft SQL Server Date Function	Notes
MONTH	
MONTHNAME	SQL Server 2014 only.
QUARTER	SQL Server 2014 only.
SECOND	SQL Server 2014 only.
WEEK	SQL Server 2014 only.
YEAR	

Microsoft SQL Server Encryption Function Support

TDV supports the encryption functions listed in the table below for Microsoft SQL Server 2008. If the SQL Server string data type is not CHAR or VARCHAR, the results are different when function is not pushed.

Microsoft SQL Server Encryption Function	Notes
HASHMD2	
HASHMD5	
HASHSHA	
HASHSHA1	

Microsoft SQL Server Numeric Function Support

TDV supports the numeric functions listed in the table below for Microsoft SQL Server.

Notes:

- These numeric functions do not support unique identifiers.

- Microsoft SQL Server 2005, 2008, and 2012 support floating point and numeric data types for both arguments of a modulo (%) operator.

Microsoft SQL Server Numeric Function	Notes
ABS	
ACOS	
ASIN	
ATAN	
CEILING	
COS	
COT	
DEGREES	
EXP	
FLOOR	
LOG	
PI	
POWER	
RADIANS	
ROUND	
SIN	
SQRT	
TAN	

Microsoft SQL Server Time Function Support

TDV supports the time function listed in the table below for Microsoft SQL Server.

Function	Notes
EXTRACT	

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>
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.