



# TIBCO® Data Virtualization

## Vertica Adapter Guide

Version 8.7.0 | October 2023

# Contents

---

<b>Contents</b>	<b>2</b>
<b>TDV Vertica Adapter</b>	<b>4</b>
Introduction	4
Obtain and Install the Driver for Vertica	4
Datasource Configuration	5
Basic Tab	5
Advanced Tab	10
Using Pass-Through Introspection with Vertica Data Source Clients	26
Vertica Data Source Limitations	26
Vertica Caching Limitations	26
Data Type Mappings	27
Vertica to TDV Data Types	27
Vertica Cache Mapping	29
Vertica Function Support	31
Vertica Aggregate Function Support	32
Vertica Analytic Function Support	34
Vertica Binary Function Support	35
Vertica Character Function Support	36
Vertica Conditional Function Support	39
Vertica Conversion Function Support	40
Vertica Date Function Support	40
Vertica Numeric Function Support	44
Vertica OLAP Analytic Function Support	45
Vertica Specific Properties	47
Vertica Time Series Function Support	48
References	49

<b>TIBCO Product Documentation and Support Services .....</b>	<b>51</b>
How to Access TIBCO Documentation .....	51
How to Contact TIBCO Support .....	52
Release Version Support .....	52
How to Join TIBCO Community .....	53
<b>Legal and Third-Party Notices .....</b>	<b>54</b>

# TDV Vertica Adapter

---

## Introduction

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

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

[Datasource Configuration](#)

[Using Pass-Through Introspection with Vertica Data Source Clients](#)

[Vertica Data Source Limitations](#)

[Data Type Mappings](#)

[Vertica Function Support](#)

[References](#)

## Obtain and Install the Driver for Vertica

Using Vertica with TDV requires that you obtain and install the Vertica JDBC driver according to instructions from Vertica.

For instructions on installing the Vertica JDBC drivers, see Client driver install procedures in the Vertica Programmer's Guide.

### To obtain and install JDBC drivers for HP Vertica

Install the driver files according to the instructions from Vertica.

- For Vertica 5.0, use the Vertica 4.1.19 driver (vertica\_4.1.19\_jdk\_5.jar).
- For Vertica 6.1, use the Vertica 6.1.2 driver (vertica-jdk5-6.1.2-0.jar).

1. Paste the driver files to the following TDV installation directory:

```
<TDV_install_dir>\conf\adapters\system\vertica_<versionx_y>
```

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	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 to enable Kerberos security through keytab files. Type the full path to the keytab file
Service 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

""

## 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 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 Kerberos authentication.

## Service Principal Name

Indicates the Kerberos Service Principal Name.

## Data Type

String

## Default Value

""

## Remarks

This field is available only if you choose 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.

<a href="#">Connection URL Pattern</a>	A template for generating a URL to connect to the physical data source.
<a href="#">Connection URL String</a>	The URL string generated from the connection URL pattern with the connection information you provide.
<a href="#">JDBC Connection Properties</a>	Lets you specify property-value pairs to pass to the JDBC data source
<a href="#">Connection Pool Maximum Size</a>	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.
<a href="#">Connection Pool Minimum Size</a>	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.

<a href="#">Enable Native Data Loading</a>	Lets the data source use its proprietary functionality to optimize performance.
<a href="#">Collation Sensitive</a>	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="#">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="#">Enable Bulk Import/Exportd</a>	Setting this option indicates that you want to use Vertica's Bulk Load utility to speed up your query times. For a given query, cardinality information is used to decide whether to use Bulk Load or JDBC default loading.

<a href="#">Enable Export To Another Vertica Database</a>	When set, allows data to be exported to another Vertica database. You need to name each database that is available to accept the exported data. TDV uses the CONNECT and EXPORT commands to establish the connections between the data ship source and the data ship target.
<a href="#">Exported Database</a>	Name of the Vertica database to which you want to export data.
Path of data source	The TDV full pathname to the other Vertica data source. For example: <pre>/shared/vertica</pre>

## 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 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

Setting this option indicates that you want to use Vertica's Bulk Load utility to speed up your query times. For a given query, cardinality information is used to decide whether to use Bulk Load or JDBC default loading.

## Data Type

Bool

## Default Value

True

## Remarks

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

## Enable Export To Another Vertica Database

When set, allows data to be exported to another Vertica database. You need to name each database that is available to accept the exported data. TDV uses the CONNECT and EXPORT commands to establish the connections between the data ship source and the data ship target.

## Data Type

Bool

## Default Value

False

## Remarks

To make changes in this field, Is dataship target 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.

## Exported Database

Name of the Vertica database to which you want to export data.

## Data Type

String

## Default Value

""

## Remarks

None

# Using Pass-Through Introspection with Vertica Data Source Clients

When Introspect as Pass-Through Following DDL is enabled on a Vertica data source, tables created in the data source using the DDL feature are created using the same pass-through user that submitted the DDL. Enabling this will make introspection ignore the saved credential if the DDL was issued by a pass-through user.

## To set pass-through optimization for Vertica data source clients

Log into Studio as the admin user.

1. From the Administration menu, choose Configuration.
2. In the tree pane, navigate to Configuration > Data Sources > Vertica Sources.
3. Select true, to enable introspection as a pass-through user.
4. Click Apply.
5. Click OK.
6. Restart the TDV Server to implement any changes you have made.

## Vertica Data Source Limitations

This section lists the general limitations related to Vertica. Other limitations depend on the type of features you want to use with Vertica data sources.

- Vertica has no user-defined indexes. Studio cannot index Vertica data sources.
- You cannot use Boolean expressions in a projection. Rewrite the view or use a packaged query.
- TDV supports hints only after the SELECT keyword.

## Vertica Caching Limitations

Because of Vertica length limits, mapping of any data type (BINARY, CHAR, VARCHAR, BLOB, and so on) to Vertica cache with length greater than 65000 results in an error, regardless of the data source.

Similarly, Vertica only supports precision up to 15 digits. Any data that you have past 16 digits will get rounded by Vertica. This precision limitation is particularly noticeable when working with REAL, FLOAT, and DOUBLE data types.

## Data Type Mappings

### Vertica to TDV Data Types

Mapped Vertica data types have the following restrictions:

- Maximum BINARY length is 65000.
- Maximum VARBINARY length is 65000.
- Maximum CHAR length is 65000.
- Maximum VARCHAR length is 65000.

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

Vertica Data Type	TDV Data Type
BIGINT	BIGINT
BINARY	BINARY
BINARY VARYING	VARBINARY
BOOL	CHAR
BOOLEAN	BOOLEAN (See User Guide, Section Mapping of Native to TDV Data Types Across TDV Versions).
BYTEA	VARBINARY
CHAR	CHAR

Vertica Data Type	TDV Data Type
CHARACTER	CHAR
CHARACTER VARYING	VARCHAR
DATE	DATE
DATETIME	TIMESTAMP
DECIMAL	DECIMAL
DOUBLE PRECISION	DOUBLE
FLOAT	DOUBLE
INT	BIGINT
INTEGER	BIGINT
INTERVAL	VARCHAR
MONEY	DECIMAL
NUMBER	NUMBER
NUMERIC	DECIMAL
RAW	VARBINARY
REAL	DOUBLE
SMALLDATETIME	TIMESTAMP
SMALLINT	BIGINT
TIME	TIME
TIMESTAMP	TIMESTAMP

Vertica Data Type	TDV Data Type
TINYINT	BIGINT
VARBINARY	VARBINARY
VARCHAR	VARCHAR

## Vertica Cache Mapping

Because of Vertica length limits, mapping of any data type (BINARY, CHAR, VARCHAR, BLOB, and so on) to Vertica cache with length greater than 65000 results in an error.

The data type mappings for caches stored on Vertica are as follows.

Data Type	Native Type
BINARY(n)	BINARY(n)
BIGINT	INT8
BIT	BOOLEAN
BLOB	VARBINARY(n)
BOOL	BOOLEAN
BOOLEAN	BOOLEAN
CHAR	CHAR(n)
CLOB	VARBINARY(n)
DATE	DATE
DECIMAL	DECIMAL(p,s)

Data Type	Native Type
DOUBLE	DOUBLE PRECISION
FLOAT	FLOAT
INTEGER	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)
LONG	BIGINT
NCLOB	VARBINARY(n)

Data Type	Native Type
NUMERIC	NUMERIC(p,s)
REAL	REAL
SMALLINT	SMALLINT
TIME	TIME, TIMEZ
TIMESTAMP	TIMESTAMP
TINYINT	TINYINT
VARBINARY(n)	VARBINARY(n)
VARCHAR	VARCHAR(n)

## Vertica Function Support

TDV supports the following types of functions for Vertica:

- [Vertica Aggregate Function Support](#)
- [Vertica Analytic Function Support](#)
- [Vertica Binary Function Support](#)
- [Vertica Character Function Support](#)
- [Vertica Conditional Function Support](#)
- [Vertica Conversion Function Support](#)
- [Vertica Date Function Support](#)
- [Vertica Numeric Function Support](#)
- [Vertica OLAP Analytic Function Support](#)
- [Vertica Time Series Function Support.](#)

All aggregate, date/time, formatting, and string functions are supported in pass-through mode.

In Vertica 6.1, TDV supports:

- EXCEPT operators
- INTERSECT operators
- WITH clause, with or without column aliasing

In Vertica 5.0 or 6.1, TDV supports:

- Queries with a WHERE filter on date, time, and timestamp columns
- Subqueries in EXISTS clauses
- Subqueries in IN clauses

TDV also supports the LIMIT clause in Vertica, but maps it to TDV syntax. For example:

```
SELECT * FROM tableA LIMIT 3 OFFSET 2
```

This is implemented with the syntax:

```
SELECT * FROM tableA OFFSET 2 ROWS FETCH 3 ROWS ONLY
```

## Vertica Aggregate Function Support

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

Vertica Aggregate Function	Notes
AVG	
CORR	
COUNT	
COVAR_POP	
COVAR_SAMP	
MAX	



Vertica Aggregate Function	Notes
MIN	
PERCENTILE_CONT	
PERCENTILE_DISC	
REGR_AVGX	
REGR_AVGY	
REGR_COUNT	
REGR_INTERCEPT	
REGR_R2	
REGR_SLOPE	
REGR_SXX	
REGR_SXY	
REGR_SYY	
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
SUM	
SUM_FLOAT	
VAR_POP	
VAR_SAMP	

Vertica Aggregate Function	Notes
VARIANCE	
VARIANCE_POP	
VARIANCE_SAMP	

## Vertica Analytic Function Support

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

Vertica Analytic Function	Notes
AVG	
CORR	
COUNT	
COVAR_POP	
COVAR_SAMP	
CUME_DIST	
DENSE_RANK	
EXPONENTIAL_MOVING_AVERAGE	.
FIRST_VALUE	
LAG	
LAST_VALUE	
LEAD	

Vertica Analytic Function	Notes
MAX	
MEDIAN	
MIN	
NTILE	
PERCENT_RANK	
PERCENTILE_CONT	
PERCENTILE_DISC	
RANK	
ROW_NUMBER	
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
SUM	
VAR_POP	
VAR_SAMP	
VARIANCE	

## Vertica Binary Function Support

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

Vertica Binary Function	Notes
INT1AND, INT2AND, INT4AND, INT8AND	
INT1NOT, INT2NOT, INT4NOT, INT8NOT	
INT1OR, INT2OR, INT4OR, INT8OR	
INT1SHL, INT2SHL, INT4SHL, INT8SHL	
INT1SHR, INT2SHR, INT4SHR, INT8SHR	
INT1XOR, INT2XOR, INT4XOR, INT8XOR	

## Vertica Character Function Support

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

Vertica Character Function	Notes
ASCII	
BITCOUNT	
BITSTRING_TO_BINARY	
BIT_AND	
BIT_LENGTH	
BIT_OR	
BIT_XOR	
BTRIM	
CHAR_LENGTH	

Vertica Character Function	Notes
CHARACTER_LENGTH	
CHR	
CONCAT	
GREATEST	
NEX_TO_BINARY	
INET_ATON	
INET_NTOA	
INITCAP	
INSERT	String-manipulation function.
INSTR	
ISUTF8	
LEAST	
LEFT	
LENGTH	
LOWER	
LPAD	
LTRIM	
MD5	
OCTET_LENGTH	

Vertica Character Function	Notes
OVERLAYB	
POSITION	
QUOTE_IDENT	
QUOTE_LITERAL	
REGEXP_REPLACE	
REPEAT	
REPLACE	
RIGHT	
RPAD	
RTRIM	
SPACE	
SPLIT_PART	
STRPOS	
SUBSTR	
SUBSTRING	
TRANSLATE	
TRIM	
TRUNC	
UPPER	

Vertica Character Function	Notes
V6_ATON	
V6_NTOA	
V6_SUBNETA	
V6_SUBNETN	
V6_TYPE	

## Vertica Conditional Function Support

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

Vertica Conditional Function	Notes
COALESCE	
CONDITIONAL_CHANGE_EVENT	
CONDITIONAL_TRUE_EVENT	
DECODE	
IFNULL	
ISNULL	
NULLIF	
NVL	
NVL2	

## Vertica Conversion Function Support

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

Vertica Conversion Function	Notes
CAST	
TO_BITSTRING	
TO_CHAR	
TO_DATE	
TO_HEX	
TO_NUMBER	
TO_TIMESTAMP	
TO_TIMESTAMP_TZ	Vertica 6.1 only.

## Vertica Date Function Support

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

With Vertica date functions, push results may differ from no-push results. For example:

- Although Vertica accepts `TIMESTAMP` as an argument for `MONTHS_BETWEEN`, it ignores the time part while calculating. TDV honors the time part while calculating.
- `DATEDIFF` calculates results according to ticks (boundaries) crossed within a date or time range (counting the enddate but not the startdate). For years, the boundary is January 1. Months are based on calendar months, not the days within months. Weeks start at midnight on Sunday. Days are based on calendar days, not the hours within days, and so on.



Vertica Date Function	Notes
ADD_MONTHS	
AT TIME ZONE	
CLOCK_TIMESTAMP	
CURRENT_DATE	
CURRENT_TIME	
CURRENT_TIMESTAMP	
DATEDIFF	
DATE_PART	
DATE_TRUNC	
DAY	
DAYOFMONTH	
DAYOFWEEK	
DAYOFWEEK_ISO	
DAYOFYEAR	
DAYS	
EXTRACT	.
GETDATE	Not supported. Use SYSDATE (identical).
GETUTCDATE	

Vertica Date Function	Notes
HOUR	
ISFINITE	
JULIAN_DAY	
LAST_DAY	
LOCALTIME	
LOCALTIMESTAMP	
MICROSECOND	
MIDNIGHT_SECONDS	
MINUTE	
MONTH	
MONTHS_BETWEEN	Because
NEW_TIME	Vertica 6.1 only.
NEXT_DAY	
NOW	
QUARTER	
ROUND	For date/time arguments.
SECOND	
STATEMENT_TIMESTAMP	
SYSDATE	

Vertica Date Function	Notes
TIMEOFDAY	
TIMESTAMPADD	<p>Keywords that TDV requires in the TDV source code as the first argument (no quotation marks):</p> <ul style="list-style-type: none"> <li>• SQL_TSI_FRAC_SECOND</li> <li>• SQL_TSI_SECOND</li> <li>• SQL_TSI_MINUTE</li> <li>• SQL_TSI_HOUR</li> <li>• SQL_TSI_DAY</li> <li>• SQL_TSI_WEEK</li> <li>• SQL_TSI_MONTH</li> <li>• SQL_TSI_QUARTER</li> <li>• SQL_TSI_YEAR</li> </ul>
TIMESTAMP_ROUND	
TIMESTAMP_TRUNC	
TIMESTAMPDIFF	<p>Keywords that TDV requires in the TDV source code as the first argument (no quotation marks):</p> <ul style="list-style-type: none"> <li>• SQL_TSI_FRAC_SECOND</li> <li>• SQL_TSI_SECOND</li> <li>• SQL_TSI_MINUTE</li> <li>• SQL_TSI_HOUR</li> <li>• SQL_TSI_DAY</li> <li>• SQL_TSI_WEEK</li> <li>• SQL_TSI_MONTH</li> </ul>

Vertica Date Function	Notes
	<ul style="list-style-type: none"> <li>SQL_TSI_QUARTER</li> <li>SQL_TSI_YEAR</li> </ul>
TRUNC	
WEEK	
WEEK_ISO	
YEAR	
YEAR_ISO	

## Vertica Numeric Function Support

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

Vertica Numeric Function	Notes
ABS	
ACOS	
ASIN	
ATAN	
ATAN2	
CBRT	
CEILING	
COS	

Vertica Numeric Function	Notes
COT	
DEGREES	
EXP	
FLOOR	
LN	
LOG	
MOD	
PI	
POWER	
RADIANS	
RANDOM	
ROUND	
SIGN	
SIN	
SQRT	
TAN	

## Vertica OLAP Analytic Function Support

TDV supports the Vertica OLAP analytic functions shown in the table below. These analytic functions are supported in pass-through mode.

Each of these functions returns the same number of rows as the input. These functions operate on groups of rows defined by frame clauses and window partitioning rather than by a GROUP BY clause.

Vertica OLAP Analytic Function	Notes
AVG	
COUNT	
CUME_DIST	
DENSE_RANK	
EXPONENTIAL_MOVING_AVERAGE	
FIRST_VALUE	
LAG	
LAST_VALUE	
LEAD	
MAX	
MEDIAN	
MIN	
NTILE	
PERCENT_RANK	
PERCENTILE_CONT	
PERCENTILE_DISC	
RANK	

Vertica OLAP Analytic Function	Notes
ROW_NUMBER	
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
SUM	
VAR_POP	
VAR_SAMP	
VARIANCE	

## Vertica Specific Properties

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

Port	Port number for the data source to connect with the host.  The Port number for Vertica is 5433.
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.
Enable Bulk Export/Load	Setting this option indicates that you want to use Vertica's Bulk Load utility to speed up your query times. For a given query, cardinality information is used to decide whether to use Bulk Load or JDBC default loading.
Enable Export To Another Vertica Database	When set, allows data to be exported to another Vertica database. You need to name each database that is available to accept the exported data. TDV uses the CONNECT and EXPORT commands to establish the connections between the data ship source and the data ship target.
Exported database name	Name of the Vertica database to which you want to export data.
Path of data source	TheTDV full pathname to the other Vertica data source. For example: <pre>/shared/vertica</pre>

## Vertica Time Series Function Support

TDV supports the Vertica time series functions listed in the table below, which are used in TIMESERIES clauses. For a complete description of the functions, see the Vertica Analytic



Database SQL Reference Manual. For details on how to use them with TDV, see “Data Ship Performance Optimization” in the TDV User Guide.

Vertica Time Series Function	Notes
TIME_SLICE	
TIMESERIES	
TO_FIRST_VALUE	
TO_LAST_VALUE	
TS_FIRST_VALUE	
TS_LAST_VALUE	

## 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.