



TIBCO® Data Virtualization

Greenplum Adapter Guide

Version 8.7.0 | October 2023

Contents

Contents	2
TDV Greenplum Adapter	4
Introduction	4
Datasource Configuration	4
Basic Tab	4
Advanced Tab	11
Greenplum to TDV Data Types	25
Greenplum Cache Mapping	28
Greenplum Function Support	30
Greenplum Aggregate Function Support	31
Greenplum Analytic Function Support	31
Greenplum Analytic Aggregate Function Support	33
Greenplum Binary Function Support	35
Greenplum Character Function Support	36
Greenplum Conditional Function Support	39
Greenplum Conversion Function Support	40
Greenplum Date Function Support	40
Greenplum Numeric Function Support	43
Greenplum Time Function Support	45
References	45
TIBCO Product Documentation and Support Services	47
How to Access TIBCO Documentation	47
How to Contact TIBCO Support	48
Release Version Support	48
How to Join TIBCO Community	49

Legal and Third-Party Notices	50
--	-----------

TDV Greenplum Adapter

Introduction

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

[Datasource Configuration](#)

[Greenplum to TDV Data Types](#)

[Greenplum 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.
Port	Port number for the data source to connect with the host.

	The Port number for Greenplum is 5432.
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	Indicates the Authentication type. Choose between BASIC and KERBEROS.
Kerbos Server Name	Name of the Kerberos server.
Include Realm	This field is enabled only when KERBEROS Authentication is used. If this field is disabled, the Greenplum Database role name is the Kerberos principal name without the Kerberos realm. If it is Enabled, the Kerberos realm is not stripped from the Greenplum Database rolename.
Keytab File	Use to enable Kerberos security through Keytab files. Type the full path to the Keytab file.
Service Principal Name	This field is available only if you choose Kerberos authentication.

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

Authentication

Indicates the Authentication type. Choose between BASIC and KERBEROS.

Data Type

String

Default Value

BASIC

Remarks

None

Kerbos Server Name

Name of the Kerberos server.

Data Type

String

Default Value

""

Remarks

None

Include Realm

This field is enabled only when KERBEROS Authentication is used. If this field is disabled, the Greenplum Database role name is the Kerberos principal name without the Kerberos realm. If it is Enabled, the Kerberos realm is not stripped from the Greenplum Database rolename.

Data Type

Bool

Default Value

False

Remarks

None

Keytab File

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

Data Type

String

Default Value

""

Reamrks

None

Service Principal Name

This field is available only if you choose Kerberos authentication.

Data Type

This field is available only if you choose Kerberos authentication.

Default Value

""

Remarks

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.
Advanced Tab	Lets the data source use its proprietary functionality to optimize performance.
Collation Sensitive	TDV does not use the SORT MERGE join algorithm if any data source involved in the join is marked Collation Sensitive.
Concurrent Request Limit	Works with the Massively Parallel Processing engine configuration parameters to control the amount of parallelization for the queries for a particular data source.
Is dataship source	This must be checked if the physical data source might be used as a source of shipped tables to another data

	ship enabled data source. Check Is dataship source for all data sources so that the TDV Server can analyze the query and determine the side that would be best to ship to based on expected or estimated query node cardinality.
Is dataship target	This must be checked if the physical data source might be used to receive shipped tables from another data ship enabled data source. Check Is dataship target for all data sources so that the TDV Server can analyze the query and determine the side that would be best to ship to based on expected or estimated query node cardinality.
Lower bound for dataship/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. Required for DB2, make sure that this name matches a schema name known to TDV. Case should match exactly.
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.

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

This must be checked if the physical data source might be used as a source of shipped tables to another data ship enabled data source. Check Is dataship source for all data sources so that the TDV Server can analyze the query and determine the side that would be best to ship to based on expected or estimated query node cardinality.

Data Type

Bool

Default Value

False

Remarks

None

Is dataship target

This must be checked if the physical data source might be used to receive shipped tables from another data ship enabled data source. Check Is dataship target for all data sources so that the TDV Server can analyze the query and determine the side that would be best to ship to based on expected or estimated query node cardinality.

Data Type

Bool

Default Value

False

Remarks

Is dataship source should be enabled in order to make changes to this field.

Lower bound for dataship/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

Is dataship source should be enabled in order to make changes to this field.

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.

Required for DB2, make sure that this name matches a schema name known to TDV. Case should match exactly.

Data Type

String

Default Value

None

Remarks

None

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

None

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.

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.

Data Type

Bool

Default Value

True

Remarks

Is dataship source should be enabled in order to make changes to this field.

Greenplum to TDV Data Types

This section provides the data type mappings from Greenplum to TDV data types.

Unsupported Data Types

Functions are not supported for operations on the following data types, which are mapped but not verified by TDV: CID, CIDR, INET, LINE, LSEG, MACADDR, PATH, POINT, POLYGON. For example, POINT should have a format like number,number. If a value with another format is inserted, an exception is thrown.

Type Promotion in Greenplum

In some circumstances, Greenplum performs type promotion that causes results to differ between push and no-push query execution. For example, with arithmetic operators a FLOAT4 column is converted to a FLOAT8/DOUBLE data type, and the Greenplum results have extra digits in the mantissa.

Data Type Mapping

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

Greenplum Data Type	TDV Data Type
BIGINT	BIGINT
BIGSERIAL	BIGINT
BIT	CHAR
BOOL	BOOLEAN (See Mapping of Native to TDV Data Types Across TDV Versions, page 555.)
BOOLEAN	BOOLEAN (See Mapping of Native to TDV Data Types Across TDV Versions, page 555.)
BOX	VARCHAR
BYTEA	BLOB
CHAR	CHAR
CHARACTER	CHAR
CHARACTER VARYING	VARCHAR
CID	CHAR
CIDR	VARCHAR
CIRCLE	VARCHAR
DATE	DATE
DECIMAL	DECIMAL
DOUBLE PRECISION	DOUBLE

Greenplum Data Type	TDV Data Type
FLOAT4	FLOAT
FLOAT8	DOUBLE
INET	VARCHAR
INT2	SMALLINT
INT4	INTEGER
INT8	BIGINT
INTEGER	INTEGER
INTERVAL	VARCHAR
LINE	VARCHAR
LSEG	VARCHAR
MACADDR	VARCHAR
MONEY	DECIMAL
NUMERIC	Arbitrary Numeric
OID	BLOB
PATH	VARCHAR
POINT	CHAR
POLYGON	VARCHAR
REAL	REAL
SERIAL	INTEGER

Greenplum Data Type	TDV Data Type
SMALLINT	SMALLINT
TEXT	CLOB
TIME	TIME
TIMESTAMP	TIMESTAMP
UUID	CHAR
VARBIT	VARCHAR
XID	INTEGER
XML	XML

Greenplum Cache Mapping

The data type mappings for caches stored on Greenplum areas follows.

Data Type	Native Type
BIGINT	BIGINT
BINARY(n)	BYTEA OID
BIT	SMALLINT
BLOB	OID
BOOLEAN	BOOLEAN
CHAR(n)	CHAR(n) TEXT

Data Type	Native Type
CLOB	TEXT
DATE	DATE
DECIMAL	DECIMAL(p,s)
DOUBLE	DOUBLE PRECISION
FLOAT	REAL
INTEGER	INTEGER
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)

Data Type	Native Type
NUMERIC	NUMERIC(p,s)
REAL	REAL
SMALLINT	SMALLINT
TIME	TIME
TIMESTAMP	VARCHAR(26)
TINYINT	SMALLINT
VARBINARY(n)	BYTEA OID
VARCHAR(n)	VARCHAR(n) TEXT
XML	XML

Greenplum Function Support

Greenplum Database is based on PostgreSQL and adheres to the same SQL structure and syntax (with minor exceptions).

TDV supports the following types of functions for Greenplum:

- [Greenplum Aggregate Function Support](#)
- [Greenplum Analytic Function Support](#)
- [Greenplum Analytic Aggregate Function Support](#)
- [Greenplum Binary Function Support](#)
- [Greenplum Character Function Support](#)
- [Greenplum Conditional Function Support](#)
- [Greenplum Conversion Function Support](#)

- [Greenplum Date Function Support](#)
- [Greenplum Numeric Function Support](#)
- [Greenplum Time Function Support](#)

Greenplum Aggregate Function Support

TDV supports the aggregate functions listed in the table below for Greenplum. DISTINCT is supported for all of these functions.

Greenplum Aggregate Function	Notes
AVG	Push supported.
BIT_AND	Push supported.
BIT_OR	Push supported.
COUNT	Push supported.
MAX	Push supported.
MIN	Push supported.
SUM	Push supported.

Greenplum Analytic Function Support

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

The following functions can not be pushed:

- EXP_WEIGHTED_AVG
- FIRST_VALUE_IGNORE_NULLS
- LAST_VALUE_IGNORE_NULLS
- NTH_VALUE_FROM_LAST

- NTH_VALUE_FROM_LAST_IGNORE_NULLS
- NTH_VALUE
- NTH_VALUE_IGNORE_NULLS
- RATIO_TO_REPORT
- TIMESERIES

Greenplum Analytic Function	Notes
AVG	Push supported.
CORR	Push supported.
COUNT	Push supported.
COVAR_POP	Push supported.
COVAR_SAMP	Push supported.
CUME_DIST	Push supported.
DENSE_RANK	Push supported.
FIRST_VALUE	Push supported.
LAG	Push supported.
LAST_VALUE	Push supported.
LEAD	Push supported.
MAX	Push supported.
MIN	Push supported.
NTILE	Push supported.
PERCENT_RANK	Push supported.

Greenplum Analytic Function	Notes
RANK	Push supported.
ROW_NUMBER	
STDDEV	DISTINCT supported. Push supported. TDV's implementation of STDDEV upcasts 32 bit float to 64 bit double. The result is a double
STDDEV_POP	Push supported.
STDDEV_SAMP	Push supported.
VAR_POP	Push supported.
VAR_SAMP	Push supported.
VARIANCE	DISTINCT supported. Push supported.
VARIANCE_POP	Push supported.
VARIANCE_SAMP	Push supported.

Greenplum Analytic Aggregate Function Support

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

The following functions can not be pushed:

- CORR_SPEARMAN
- LISTAGG
- MEDIAN
- PERCENTILE_CONT

- PERCENTILE_DISC
- XMLAGG

Greenplum Analytic Aggregate Function	Notes
CORR	Push supported.
COVAR_POP	Push supported.
COVAR_SAMP	Push supported.
REGR_AVGX	Push supported.
REGR_AVGY	Push supported.
REGR_COUNT	Push supported.
REGR_INTERCEPT	Push supported.
REGR_R2	Push supported.
REGR_SLOPE	Push supported.
REGR_SXX	Push supported.
REGR_SXY	Push supported.
REGR_SYY	Push supported.
STDDEV	Push supported.
STDDEV_POP	Push supported.
STDDEV_SAMP	Push supported.
VAR_POP	Push supported.

Greenplum Analytic Aggregate Function	Notes
VAR_SAMP	Push supported.
VARIANCE	Push supported.
VARIANCE_POP	Push supported.
VARIANCE_SAMP	Push supported.

Greenplum Binary Function Support

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

Greenplum Binary Function	Notes
INT1AND	Push supported.
INT2AND	Push supported.
INT4AND	Push supported.
INT8AND	Push supported.
INT1OR	Push supported.
INT2OR	Push supported.
INT4OR	Push supported.
INT8OR	Push supported.
INT1SHL	Push supported.
INT2SHL	Push supported.

Greenplum Binary Function	Notes
INT4SHL	Push supported.
INT8SHL	Push supported.
INT1SHR	Push supported.
INT2SHR	Push supported.
INT4SHR	Push supported.
INT8SHR	Push supported.
INT1XOR	Push supported.
INT2XOR	Push supported.
INT4XOR	Push supported.
INT8XOR	Push supported.
INT1NOT	Push supported.
INT2NOT	Push supported.
INT4NOT	Push supported.
INT8NOT	Push supported.

Greenplum Character Function Support

TDV supports the character functions listed in the table below for Greenplum

The following functions can not be pushed:

- DLE_DST
- INSERT
- LE_DST

- LOCATE
- PARTIAL_STRING_MASK

Greenplum Character Function	Notes
ASCII	Push supported.
BIT_LENGTH	Push supported.
BTRIM	Push supported.
CHAR_LENGTH	Push supported.
CHARACTER_LENGTH	Push supported.
CHR	Push supported.
CONCAT	Results might differ between pushed and not pushed, even if the Ignore Trailing Space setting of the Greenplum data source is the same as that of TDV, because the Greenplum database always trims trailing spaces.
FIND	Push supported.
INITCAP	Push supported.
INSTR	Push supported.
LCASE	Push supported.
LENGTH	Push supported.
LOWER	Push supported.
LPAD	Push supported.

Greenplum Character Function	Notes
LTRIM	Push supported.
POSITION	Results might differ between pushed and not pushed, even if the Ignore Trailing Space setting of the Greenplum data source is the same as that of TDV, because the Greenplum database always trims trailing spaces.
REPEAT	Push supported.
REPLACE	Results might differ between pushed and not pushed, even if the Ignore Trailing Space setting of the Greenplum data source is the same as that of TDV, because the Greenplum database always trims trailing spaces.
RPAD	Push supported.
RTRIM	Push supported.
SPACE	Push supported.
STRPOS	Push supported.
SUBSTR	Push supported.
SUBSTRING	Push supported.
TO_HEX	Push supported.
TRANSLATE	Push supported.
TRIM	Push supported.
TRIM(LEADING FROM)	Push supported.
TRIM(TRAILING FROM)	Push supported.

Greenplum Character Function	Notes
UCASE	Push supported.
UNICHR	Push supported.
UNICODE	Push supported.
UPPER	Push supported.

Greenplum Conditional Function Support

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

The following functions can not be pushed:

- DECODE
- ISNULL
- ISNUMERIC

Greenplum Conditional Function	Notes
COALESCE	Push supported.
GREATEST	Push supported.
IFNULL	Push supported.
LEAST	Push supported.
NULLIF	Push supported.
NVL	Push supported.
NVL2	Push supported.

Greenplum Conversion Function Support

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

The following functions can not be pushed:

- PARSE_DATE
- PARSE_TIME
- TIMESTAMP
- TO_TIMESTAMP_TZ

Greenplum Conversion Function	Notes
CAST	Push supported.
PARSE_TIMESTAMP	Push supported.
TO_CHAR	Push supported.
TO_DATE	Push supported.
TO_NUMBER	Push supported.
TO_TIMESTAMP	Push supported.

Greenplum Date Function Support

TDV supports the date functions listed in the table below for Greenplum

The following functions can not be pushed:

- DATEADD
- DATENAME
- DAYNAME
- DAYS_BETWEEN
- DBTIMEZONE

- LAST_DAY
- MONTHS_BETWEEN
- NEW_TIME
- NEXT_DAY
- NUMTODSINTERVAL
- NUMTOYMINTERVAL
- TIME_SLICE
- TRUNC
- TZCONVERTOR
- UTC_TO_TIMESTAMP

Greenplum Date Function	Notes
ADD_MONTHS	Push supported.
CLOCK_TIMESTAMP	Push supported.
CURRENT_DATE	Push supported.
CURRENT_TIME	Push supported.
CURRENT_TIMESTAMP	Push supported.
DATE_ADD	Push supported.
DATE_PART	Push supported.
DATE_SUB	Push supported.
DATE_TRUNC	Push supported.
DATEPART	Push supported.

Greenplum Date Function	Notes
DATETRUNC	Push supported.
DAY	Push supported.
EXTRACT(DAY FROM)	Push supported.
EXTRACT(DOW FROM)	Push supported.
EXTRACT(DOY FROM)	Push supported.
EXTRACT(EPOCH FROM)	Push supported.
EXTRACT(HOUR FROM)	Push supported.
EXTRACT(MICROSECOND FROM)	Push supported.
EXTRACT(MILLISECOND FROM)	Push supported.
EXTRACT(MINUTE FROM)	Push supported.
EXTRACT(MONTH FROM)	Push supported.
EXTRACT(QUARTER FROM)	Push supported.
EXTRACT(SECOND FROM)	Push supported.
EXTRACT(WEEK FROM)	Push supported.
EXTRACT(YEAR FROM)	Push supported.
FORMAT_DATE	Push supported.
LOCALTIME	Push supported.
LOCALTIMESTAMP	Push supported.
MONTH	Push supported.

Greenplum Date Function	Notes
NOW	Push supported.
TIMEOFDAY	Push supported.
YEAR	Push supported.

Greenplum Numeric Function Support

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

The following functions can not be pushed:

- COSH
- SINH
- TANH
- FACTORIAL
- ROWNUM

Greenplum Numeric Function	Notes
ABS	Push supported.
ACOS	Push supported.
ASIN	Push supported.
ATAN	Push supported.
ATAN2	Push supported.
CBRT	Push supported.
CEIL	Push supported.

Greenplum Numeric Function	Notes
CEILING	Push supported.
COS	Push supported.
COT	Push supported.
DEGREES	Push supported.
EXP	Push supported.
FLOOR	Push supported.
LN	Push supported.
LOG	Push supported.
MOD	Push supported.
NUMERIC_LOG	Push supported.
PI()	Push supported.
POW	Push supported.
POWER	Push supported.
RADIANS	Push supported.
RAND	Push supported.
RANDOM	Push supported.
ROUND	Push supported.
SIGN	Push supported.
SIN	Push supported.

Greenplum Numeric Function	Notes
SQRT	Push supported.
TAN	Push supported.
TRUNC	Push supported.

Greenplum Time Function Support

TDV supports the time function listed in the table below for Greenplum.

Greenplum Time 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>

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

TIBCO Product Documentation and Support Services

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

How to Access TIBCO Documentation

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

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

Product-Specific Documentation

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

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

TDV Adapter Guides

TDV Data Source Toolkit Guide (Formerly Extensibility Guide)

- **References**

TDV Reference Guide

TDV Application Programming Interface Guide

- **Other**

TDV Business Directory Guide

TDV Discovery Guide

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

How to Contact TIBCO Support

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

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

Release Version Support

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

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

How to Join TIBCO Community

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

Legal and Third-Party Notices

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

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

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

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

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

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

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

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

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

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

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

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

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