



TIBCO® Data Virtualization

ComputeDB Adapter Guide

Version 8.7.0 | October 2023

Contents

Contents	2
ComputeDB Adapter	4
Introduction	4
Datasource Configuration	4
Basic Tab	4
Advanced Tab	8
ComputeDB Data Source Limitations	16
ComputeDB Caching Characteristics	17
Data Type Mappings	18
ComputeDB to TDV Data Types	18
ComputeDB Cache Mapping	20
ComputeDB Function Support	21
ComputeDB Aggregate Function Support	22
ComputeDB Character Function Support	22
ComputeDB Conversion Function Support	23
ComputeDB Date Function Support	24
ComputeDB Numeric Function Support	25
ComputeDB SparkSQL Function Support	26
ComputeDB Specific Properties	26
References	27
TIBCO Product Documentation and Support Services	28
How to Access TIBCO Documentation	28
How to Contact TIBCO Support	29
Release Version Support	29
How to Join TIBCO Community	30

Legal and Third-Party Notices	31
--	-----------

ComputeDB Adapter

Introduction

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

[Datasource Configuration](#)

[ComputeDB Data Source Limitations](#)

[ComputeDB Caching Characteristics](#)

[Data Type Mappings](#)

[ComputeDB 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.

Data source Name	The name of the data source.
Complete Connection URL String	A URL to connect to the physical data source.

Leader Node Port	The Lead node port used for submitting ComputeDB jobs. The default port is 8090.
Login/Password	User name and password required to access the data source. 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.
Pass-through Login	Flag to indicate whether pass-through login is enabled or not.
Authentication Type	The type of authentication used by the data source.

Data source Name

The name of the data source.

Data Type

string

Default Value

""

Complete Connection URL String

A URL to connect to the physical data source.

Data Type

String

Default Value

""

Remarks

TDV does not validate modifications. The data source adapter might not validate changes. Refer to the datasource specific connection properties section for details.

Leader Node Port

The Lead node port used for submitting ComputeDB jobs. The default port is 8090.

Data Type

Numeric

Default Value

8090

Remarks

None

Login/Password

The login name used to authenticate the data source connection,

Data Type

string

Default Value

“”

Remarks

None

Password

The password used to authenticate the data source connection,

Data Type

string

Default Value

""

Remarks

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

Following are the available Authentication types:

BASIC

KERBEROS

NTLM

NEGOTIATE

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

Advanced Tab

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

Transaction Isolation	The degree to which transactions are isolated from data modifications made by other transactions.
Connection Attributes	Lets you specify property-value pairs to pass to the 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 Checkout Procedure	A procedure that returns a valid SQL statement that can be used to initialize the connection.
Connection Checkout Timeout	Time that a connection doing a checkout can remain idle without being dropped.
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 source driver doesn't support query timeout	Select or clear the check box. If cleared, specify an Execution timeout value in seconds.
Execution Timeout	The number of seconds an execution query on the data source can run before being canceled.
Execute SELECTs in separate transactions from INSERTs and UPDATEs	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.
String Comparison Parameters	Multiple Parameters that can be set to control String manipulation by the data source.

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

Connection Attributes

Lets you specify property-value pairs to pass to the data source. For example:

```
bootPassword=key attribute
```

```
collation=collation attribute
```

```
dataEncryption=true attribute
```

```
drop=true attribute
```

```
encryptionKey=key attribute
```

```
encryptionProvider=providerName attribute
```

```
encryptionAlgorithm=algorithm attribute
```

```
failover=true attribute
```

Data Type

String

Default Value

“”

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.

Data source driver doesn't support query timeout

Select or clear the check box. If cleared, specify an Execution timeout value in seconds.

Data Type

Bool

Default Value

True

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 in separate transactions from INSERTs and UPDATEs

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.

Connection Checkout Timeout

Time that a connection doing a checkout can remain idle without being dropped.

Data Type

Numeric

Default Value

45

Remarks

None

String Comparison Parameters

Following flags can be used to manage the rules of the string comparison.

Case sensitivity mismatch between TDV and data source can be ignored	Check to ignore case mismatches.
Ignore trailing space mismatches between TDV and the data source	Check to ignore trailing spaces.
Honor trailing spaces for string comparison	Check to honor trailing spaces.
Case insensitive string comparison	Ignore case.

ComputeDB Data Source Limitations

This section lists the general limitations related to ComputeDB.

- In order to introspect an empty schema, you must check the option “Detect New Resources During Re-Introspection” in the Data Source Introspection dialogue box.
- Numeric functions `pi()` and `cot()` are not supported.
- Queries resulting in a Null scan will not be pushed down.
- All table identifiers must be specified in uppercase.
- For statistics to work as expected, modify the following files while setting up ComputeDB:
 - In the `spark-defaults.conf` file, add the line `"spark.sql.thriftServer.incrementalCollect true"`
 - In the `leads.conf` file, add the line `"-spark.driver.maxResultSize=0" - "0"` (zero) is unlimited.

Note: The above 2 files can be found in the folder `$SNAPPY_HOME/conf`
- If the Transaction Isolation property of the datasource is set as Read Committed, the following functions are not recognized:
 - `ascii()`
 - `upper()`
 - `concat()`
 - `chr()`

ComputeDB Caching Characteristics

Cache target tables in ComputeDB are created either as row-store tables (in the case of Single Table Caching) or column-store tables (in the case of Multi-table caching). Row-store tables more commonly used in relational databases. The advantage of Column-store tables is that they do not have primary keys defined and are efficient while performing analysis on large data sets.

ComputeDB uses an in-built engine to move data to the cache tables. The distributed engine that is used to perform this task uses an optimization mechanism that can process data coming from multiple datasets more efficiently. The cached data is stored in memory with an option to spill to disk if necessary, which makes queries that are run against the cached data much faster.

Configuring Native Caching Option for ComputeDB

By default, TDV is configured to use ComputeDB's Native Data Loading option and the Bulk loading mode is set to "Virtual Scans".

ComputeDB supports both Single table as well as Multi table caching. Refer to the sections ***Caching to a Single-Table Database Target*** and ***Creating a Multiple Table Cache on a Database Target*** in the User Guide for more details on how to configure.

Accessing the ComputeDB Console

The ComputeDB console helps the user to monitor the resource usage by the cluster members and jobs. To access the console, start your cluster and open `http:<localhost>:5050/dashboard/` in the web browser. Refer to the [ComputeDB documentation](#) for more information.

Data Type Mappings

ComputeDB to TDV Data Types

Mapped ComputeDB data types have the following restrictions:

- Maximum BINARY length is 255.
- Maximum VARBINARY length is 4000.
- Maximim BLOB length is 4294967296
- Maximum CHAR length is 254.
- Maximum VARCHAR length is 32672
- Maximum CLOB length is 4294967296

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

ComputeDB Data Type	TDV Data Type
BIGINT	BIGINT
BINARY	BINARY
BLOB	BLOB
BOOLEAN	BOOLEAN
BYTE	SMALLINT
CLOB	CLOB
CHAR	CHAR
DATE	DATE
DECIMAL	DECIMAL
DOUBLE	DOUBLE
FLOAT	REAL
INT	INTEGER
INTEGER	INTEGER
LONG	CLOB
NUMERIC	NUMERIC
REAL	REAL
SHORT	SMALLINT
SMALLINT	SMALLINT
STRING	CLOB

ComputeDB Data Type	TDV Data Type
TIMESTAMP	TIMESTAMP
TINYINT	TINYINT
VARBINARY	VARBINARY
VARCHAR	VARCHAR

ComputeDB Cache Mapping

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

Data Type	Native Type
BIGINT	BIGINT
BINARY	BINARY
BLOB	BLOB
BOOLEAN	BOOLEAN
SMALLINT	BYTE
CLOB	CLOB
CHAR	CHAR
DATE	DATE
DECIMAL	DECIMAL
DOUBLE	DOUBLE
REAL	FLOAT

Data Type	Native Type
INTEGER	INT
INTEGER	INTEGER
CLOB	LONG
NUMERIC	NUMERIC
REAL	REAL
SMALLINT	SHORT
SMALLINT	SMALLINT
CLOB	STRING
TIMESTAMP	TIMESTAMP
TINYINT	TINYINT
VARBINARY	VARBINARY
VARCHAR	VARCHAR

ComputeDB Function Support

TDV supports the following types of functions for ComputeDB:

- [ComputeDB Aggregate Function Support](#)
- [ComputeDB Character Function Support](#)
- [ComputeDB Conversion Function Support](#)
- [ComputeDB Date Function Support](#)
- [ComputeDB Numeric Function Support](#)
- [ComputeDB SparkSQL Function Support](#)

ComputeDB Aggregate Function Support

TDV supports the aggregate functions listed in the table below for ComputeDB

ComputeDB Aggregate Function	Notes
AVG	
COUNT	
MAX	
MIN	
SUM	

ComputeDB Character Function Support

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

ComputeDB Character Function	Notes
ASCII	
CHR	
CONCAT	
INSTR	
LENGTH	
LOWER	

ComputeDB Character Function	Notes
LPAD	
POSITION	
REPLACE	
RPAD	
RTRIM	
SPACE	
SUBSTRING	
TRIM	
UPPER	

ComputeDB Conversion Function Support

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

ComputeDB Conversion Function	Notes
CAST	
PARSE_TIMESTAMP	
TO_CHAR	
TO_NUMBER	
TO_TIMESTAMP	

ComputeDB Conversion Function	Notes
FORMAT_DATE	
PARSE_DATE	

ComputeDB Date Function Support

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

ComputeDB Date Function	Notes
CURRENT_DATE	
CURRENT_TIME	
CURRENT_TIMESTAMP	
DAY	
DAYS_BETWEEN	
DBTIMEZONE	
EXTRACT	
MONTH	
MONTH_BETWEEN	
TZ_OFFSET	
TZ_CONVERTOR	
UTC_TO_TIMESTAMP	
YEAR	

ComputeDB Numeric Function Support

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

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

ComputeDB SparkSQL Function Support

TDV supports the SparkSQL functions listed in the table below for ComputeDB.

ComputeDB SparkSQL Function	Notes
INITCAP	
MD5	
SOUNDEX	
DATE_SUB	
FACTORIAL	
FLOOR	
LEAST	
LOCATE	
REGEXP_REPLACE	
TRUNC	
UNIX_TIMESTAMP	
LAST	
FIRST	

ComputeDB Specific Properties

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

Leader Node Port

The Lead node port used for submitting ComputeDB jobs. The default port is 8090.

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.