



TIBCO® Data Virtualization

Hive Adapter Guide

Version 8.7.0 | October 2023

Contents

Contents	2
TDV Hive Adapter	4
Introduction	4
Obtain and Install Drivers for Hive	4
Datasource Configuration	6
Basic Tab	7
Advanced Tab	11
Hive Compatible Data Source Limitations	22
Apache Hive Caching Limitations	23
Data Type Mappings	23
Apache Hive to TDV Data Types	23
Cloudera Hive to TDV Data Types	24
Hive Function Support	25
Apache Hive Function Support	25
Apache Hive Aggregate Function Support	25
Apache Hive Character Function Support	25
Apache Hive Conversion Function Support	26
Apache Hive Conditional Function Support	27
Apache Hive Date Function Support	27
Cloudera Hive Function Support	28
Cloudera Hive Numeric Function Support	28
Cloudera Hive Aggregate Function Support	29
Cloudera Hive Conversion Function Support	29
Cloudera Hive Date Function Support	29
Cloudera Hive Conditional Function Support	30
Cloudera Hive Character Function Support	30
References	30

TIBCO Product Documentation and Support Services	32
How to Access TIBCO Documentation	32
How to Contact TIBCO Support	33
Release Version Support	33
How to Join TIBCO Community	34
Legal and Third-Party Notices	35

TDV Hive Adapter

Introduction

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

[Obtain and Install Drivers for Hive](#)

[Datasource Configuration](#)

[Hive Compatible Data Source Limitations](#)

[Data Type Mappings](#)

[Hive Function Support](#)

[References](#)

Obtain and Install Drivers for Hive

To obtain JDBC drivers for Hive

1. From the web, locate the Hive zip files.

Choose your version from the following download locations and download hive-
<version>-bin.tar.gz.:

- <https://archive.apache.org/dist/hive>
- <https://mvnrepository.com/artifact/org.apache.hadoop>
- For instructions on using Hive JDBC driver with Kerberos, refer <https://querysurge.zendesk.com/hc/en-us/articles/115001218863-Setting-Up-a-Hive-Connection-with-Kerberos-using-Apache-JDBC-Drivers-Windows>
- <https://streever.atlassian.net/wiki/spaces/HADOOP/pages/4390924>

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

2. Unpack the zip file.
3. Copy the following JAR files for your version

Version	JAR Files
Apache Hive 2.x	commons-cli-1.2.jar
	commons-configuration-1.6.jar
	commons-logging-1.2.jar
	hadoop-auth-2.7.3.jar
	hadoop-common-2.7.3.jar
	hadoop-hdfs-2.7.3.jar
	hive-exec-2.1.1.jar
	hive-jdbc-2.1.1.jar
	hive-service-2.1.1.jar
	htrace-core-3.1.0-incubating.jar
	httpclient-4.4.jar
	httpcore-4.4.jar
	libfb303-0.9.3.jar

Version	JAR Files
	libthrift-0.9.3.jar
	log4j-core-2.4.1.jar

4. From the Apache Hadoop web site, locate and copy the hadoop-core-1.2.1.jar file.

To install the JDBC drivers for Hive

Make sure to add the JAR files to all of your TDV clients and servers.

5. Paste the necessary driver JAR files into the following TDV installation directory (or the directory for your specific version):

```
<TDV_install_dir>\conf\adapters\system\hive2
```

6. Restart the TDV Server.

To enable Kerberos authentication:

Configure the following JDBC connection URL to authenticate the connecting user on the server Hiveserver2 with Kerberos:

```
jdbc:hive2://<host>:<port>/<dbName>;principal=<HiveServer2_kerberos_principal>;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

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 Hive is 10000.
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.
Authentication Type	The type of Authentication used by the datasource.
Keytab File	This field is available only if you choose Kerberos authentication. Use to enable Kerberos security through Keytab files. Type the full path to the Keytab file.
Keytab Principal Name	The name of the Keytab principal. This field is available only if you choose KEREROS 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.

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 it 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.

Keytab Principal Name

The name of the Keytab principal.

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.

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.
Advanced Tab	Number of seconds that a connection can remain idle without being dropped from the pool when there are more than the minimum number of connections.
Maximum Connection Lifetime	The number of minutes that a connection that was returned to the pool persists if there are more open connections than the minimum pool size.
Connection Validation Query	A data-source-specific query that the TDV query engine sends to see if the data source connection is valid. This query is executed every time a connection is checked out from the pool. Enter a query that returns quickly.
Execution Timeout	The number of seconds an execution query on the data source can run before being canceled.
Execute SELECTs Independently	Lets a SELECT statement be executed using a new connection from the connection pool, and committed immediately after completion. INSERT, UPDATE, and DELETE statements are executed using the same connection as part of the transaction.
Connection Checkout Procedure	A procedure that returns a valid SQL statement that can be used to initialize the connection.
Supports Star Schema	Check only if this data source supports very large predicates and very large cardinalities for star schema semijoins.
Max Source Side Cardinality for Semi Join	See the documentation for semijoins and the TDV Administration Guide for more information.
Min Target to Source Ratio for Semi Join	Sets the minimum target-to-source ratio of cardinality for semijoins. Refer to the TDV Administration Guide for

	more information.
Max Source Side of Semi Join To Use OR Syntax	See the documentation for semijoins and the TDV Administration Guide for more information.
Enable Native Data Loading	Lets the data source use its proprietary functionality to optimize performance.
Collation Sensitive	TDV does not use the SORT MERGE join algorithm if any data source involved in the join is marked Collation Sensitive.
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.
HDFS Service Principal Name	The name of the Hadoop file system service principal.This field is available only if you choose BASIC authentication.
Cache Table Storage Format	

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.

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

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.

Connection Checkout Timeout

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

Data Type

Numeric

Default Value

45

Remarks

None

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.

HDFS Service Principal Name

The name of the Hadoop file system service principal.

Data Type

String

Default Value

""

Remarks

This field is available only if you choose BASIC authentication.

Hive Compatible Data Source Limitations

Apache Hive is a data warehouse infrastructure that lets you query and analyze large data sets stored in Apache Hadoop files. Hive uses a simple SQL-like query language called QL to query the data. TDV supports several Hive-based database solutions. For a list and associated version numbers, see the TDV Installation and Upgrade Guide. For information about how to enable security on your Hive data sources see the TDV Administration Guide.

Because queries are distributed across several Hive nodes, queries that process only a small amount of data might finish slower than expected.

Back referencing table aliases in the ORDER BY clause is not supported by Hive. In order to achieve this, you must include the columns used by the ORDER BY clause in the SELECT clause with table alias.

Apache Hive Caching Limitations

When hive is used as cache target data source, using a different data source for cache_status/cache_tracking is recommended.

Hive does not support data types BLOB, CLOB. Hence when using Hive as a cache target, tables with these data types cannot be cached.

Data Type Mappings

Apache Hive to TDV Data Types

The following limits apply to Apache Hive data type mapping:

- The maximum length for VARCHAR is 65355.

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

Apache Hive Data Type	TDV Data Type
BOOLEAN	BOOLEAN
BINARY	BINARY
CHAR	CHAR
VARCHAR	VARCHAR

Apache Hive Data Type	TDV Data Type
DECIMAL	DECIMAL
DATE	DATE
TIMESTAMP	TIMESTAMP
INTERVAL_DAY_TIME	INTERVAL\DAY\TO\SECOND
INTERVAL_YEAR_MONTH	INTERVAL\YEAR\TO\MONTH

Cloudera Hive to TDV Data Types

The following limits apply to Cloudera Hive data type mapping:

- The maximum length for VARCHAR is 65355.

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

Cloudera Hive Data Type	TDV Data Type
BOOLEAN	BOOLEAN
BINARY	BINARY
CHAR	CHAR
VARCHAR	VARCHAR
DECIMAL	DECIMAL
DATE	DATE
TIMESTAMP	TIMESTAMP

Hive Function Support

Apache Hive Function Support

TDV supports the following types of functions for Apache Hive:

[Apache Hive Aggregate Function Support](#)

[Apache Hive Character Function Support](#)

[Apache Hive Conversion Function Support](#)

[Apache Hive Conditional Function Support](#)

[Apache Hive Date Function Support](#)

Apache Hive Aggregate Function Support

TDV supports the aggregate functions listed in the table below for Apache Hive

Apache Hive Aggregate Function	Notes
NTILE	
CAST	
DATEDIFF	

Apache Hive Character Function Support

TDV supports the character functions listed in the table below for Apache Hive.

Apache Hive Character Function	Notes
CHR	
INSTR	
LOCATE	
REPEAT	
REPLACE	
TRANSLATE	
INITCAP	
SOUNDEX	

Apache Hive Conversion Function Support

Apache Hive Conversion Function	Notes
CHR	
INSTR	
LOCATE	
REPEAT	
REPLACE	
TRANSLATE	

Apache Hive Conversion Function	Notes
INITCAP	
SOUNDEX	

TDV supports the conversion functions listed in the table below for Apache Hive.

Apache Hive Conditional Function Support

TDV supports the conditional function listed in the table below for Apache Hive.

Apache Hive Conditional Function	Notes
NVL	

Apache Hive Date Function Support

TDV supports the date functions listed in the table below for Apache Hive.

Apache Hive Date Function	Notes
QUARTER	
CURRENT_DATE	
CURRENT_TIMESTAMP	
ADD_MONTHS	
LAST_DAY	

Apache Hive Date Function	Notes
NEXT_DAY	
TRUNC	
MONTHS_BETWEEN	

Cloudera Hive Function Support

TDV supports the following types of functions for Cloudera Hive:

- [Cloudera Hive Numeric Function Support](#)
- [Cloudera Hive Aggregate Function Support](#)
- [Cloudera Hive Conversion Function Support](#)
- [Cloudera Hive Date Function Support](#)
- [Cloudera Hive Conditional Function Support](#)
- [Cloudera Hive Character Function Support](#)

Cloudera Hive Numeric Function Support

Cloudera Hive Numeric Function	Notes
GREATEST	
LEAST	
SIGN	

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

Cloudera Hive Aggregate Function Support

Cloudera Hive Aggregate Function	Notes
NTILE	

TDV supports the aggregate functions listed in the table below for Cloudera Hive.

Cloudera Hive Conversion Function Support

Cloudera Hive Conversion Function	Notes
CAST	

TDV supports the conversion function listed in the table below for Cloudera Hive.

Cloudera Hive Date Function Support

Cloudera Hive Date Function	Notes
ADD_MONTHS	
DATEDIFF	
LAST_DAY	

TDV supports the date function listed in the table below for Cloudera Hive.

Cloudera Hive Conditional Function Support

Cloudera Hive Conditional Function	Notes
NVL	

TDV supports the conditional functions listed in the table below for Cloudera Hive.

Cloudera Hive Character Function Support

Cloudera Hive Character Function	Notes
INSTR	
LOCATE	
REPEAT	
TRANSLATE	
INITCAP	

TDV supports the character functions listed in the table below for Cloudera Hive.

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.