Tibco Data Virtualization®

PostgreSQL Adapter Guide

Version 8.7.0
April 2023

Contents

TDV PostgreSQL Adapter

Introduction

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

Datasource Configuration

Data Type Mappings

PostgreSQL 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 PostgreSQL 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.

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

""

Remarks

None

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.

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 of Semi Join To Use OR Syntax	See the documentation for semijoins and the TDV Administration Guide for more information.
Min Target to Source Ratio for	Sets the minimum target-to-source ratio of cardinality

Semi Join	for semijoins. Refer to the TDV Administration Guide for more information.
Max Source Side Cardinality for Semi Join	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.
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/Upper bound for dataship	TDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.

Schema Path for Temp Tables	A relative path to set the location of the temp tables on the data source. It is the name of a schema in the data source.
Temp Table Prefix	A character string addition to temporary table names so that they are recognized if they are needed.
Enable Bulk Export/Load	Takes advantage of PostgreSQL COPY command.
Enable PostgreSql dblink	Check to enable database links to improve performance if you plan to use this data source for data caching or data ship optimization. If you check this box, add one or more database links by specifying the database link name and path of the data source for each link
Database Links	If you check the box Enable PostgreSQL dblink, add one or more database links by specifying the database link name and path of the data source for each link

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

This option is available only if Is dataship source field is 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

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

Enable Bulk Export/Load

Takes advantage of PostgreSQL COPY command.

Data Type

Bool

Default Value

True

Remarks

To make changes to this field enable Is dataship source

Enable PostgreSql dblink

Check to enable database links to improve performance if you plan to use this data source for data caching or data ship optimization. If you check this box, add one or more database links by specifying the database link name and path of the data source for each link

Data Type

Bool

Default Value

False

Remarks

None

Database Links

If you check the box Enable PostgreSQL dblink, add one or more database links by specifying the database link name and path of the data source for each link

Data Type

Bool

Default Value

False

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.

PostgreSQL to TDV Data Types

The table below shows the mapping from PostgreSQL data types to TDV data types.

PostgreSQL data conversion and comparison have these traits:

- Interval years converted to months result in a TDV data type of VARCHAR.
- Interval days converted to seconds result in a TDV data type of VARCHAR.
- Timestamps with a time zone or a local time zone result in a TDV data type of TIMESTAMP.
- While PostgreSQL honors trailing spaces in general, it ignores them when comparing CHARs. When TDV is set to honor trailing spaces, a filter on a CHAR column might return different results when executed in PostgreSQL vs. TDV.

PostgreSQL Data Type	TDV Data Type
BIGINT	BIGINT
BIGSERIAL	BIGINT
BINARY DOUBLE	DOUBLE
BINARY FLOAT	REAL
BIT	CHAR
BOOL	CHAR
BOOLEAN	BOOLEAN (See the section Mapping of Native to TDV Data Types Across TDV Versions in the User Guide.)

PostgreSQL Data Type	TDV Data Type
BOX	VARCHAR
BPCHAR	CHAR
BYTEA	BLOB
CHAR	CHAR
CHARACTER	CHAR
CHARACTER VARYING	VARCHAR
CIDR	VARCHAR
CIRCLE	VARCHAR
DATE	DATE
DATETIME	TIMESTAMP
DOUBLE PRECISION	DOUBLE
FLOAT4	REAL
FLOAT8	DOUBLE
INET	VARCHAR
INT	INTEGER
INT(2)	SMALLINT
INT(4)	INTEGER
INT(8)	BIGINT
INTEGER	INTEGER

PostgreSQL Data Type	TDV Data Type
INTERVAL	VARCHAR
LINE	VARCHAR
LONG	CLOB
LSEG	VARCHAR
MACADDR	VARCHAR
MONEY	DECIMAL
NUMBER	DECIMAL
NUMERIC	Arbitrary NUMERIC
OID	BLOB
PATH	VARCHAR
POINT	CHAR
POLYGON	VARCHAR
REAL	REAL
ROWID	VARCHAR
SERIAL	INTEGER
SMALLDATETIME	TIMESTAMP
SMALLINT	SMALLINT
TEXT	CLOB
TIME	TIME

PostgreSQL Data Type	TDV Data Type
TIMESTAMP	TIMESTAMP
TIMESTAMPTZ	TIMESTAMP
TIMETZ	TIME
TINYINT	SMALLINT
UROWID	VARCHAR
UUID	CHAR
VARBIT	VARCHAR
VARCHAR	VARCHAR
VARCHAR2	VARCHAR
XID	INTEGER
XML	XML

PostgreSQL Cache Mapping

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

Data Type	Native Type
BIGINT	BIGINT
BINARY	BYTEA [regardless of size]
BIT	BOOLEAN
BLOB	BYTEA

Data Type	Native Type
BOOLEAN	BOOLEAN
CHAR	CHAR(n) TEXT
CLOB	TEXT
DATE	DATE
DECIMAL	DECIMAL TEXT
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)

Data Type	Native Type
INTERVAL SECOND	VARCHAR(30)
INTERVAL YEAR	VARCHAR(9)
INTERVAL YEAR TO MONTH	VARCHAR(12)
NUMERIC	NUMERIC(p,s) TEXT
REAL	REAL
SMALLINT	SMALLINT
TIME	TIME
TIMESTAMP	TIMESTAMP
TINYINT	SMALLINT
VARBINARY(n)	BYTEA [regardless of size]
VARCHAR(n)	VARCHAR(n) TEXT
XML	XML

PostgreSQL Function Support

TDV supports the following types of functions for PostgreSQL:

- PostgreSQL Aggregate Function Support
- PostgreSQL Analytic Aggregate Function Support
- PostgreSQL Binary Function Support
- PostgreSQL Character Function Support
- PostgreSQL Conversion Function Support

PostgreSQL Aggregate Function Support

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

PostgreSQL Aggregate Function	Notes
AVG	
CORR	
COUNT	BLOB, CLOB, and DISTINCT not supported.
COVAR_POP	
COVAR_SAMP	
MAX	
MIN	
SUM	
VARIANCE	

PostgreSQL Analytic Aggregate Function Support

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

PostgreSQL Analytic Aggregate Function	Notes
CUME_DIST	
DENSE_RANK	
NTILE	
PERCENT_RANK	
RANK	
ROW_NUMBER	
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
VAR_POP	
VAR_SAMP	
VARIANCE_POP	
VARIANCE_SAMP	

PostgreSQL Binary Function Support

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

PostgreSQL Binary Function	Notes
INT1AND, INT2AND, INT4AND, INT8AND	

PostgreSQL Binary Function	Notes
INT1NOT, INT2NOT, INT4NOT, INT8NOT	
INT1OR, INT2OR, INT4OR, INT8OR	
INT1XOR, INT2XOR, INT4XOR, INT8XOR	

PostgreSQL Character Function Support

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

PostgreSQL Character Function	Notes
ASCII	
BTRIM	
CHR	
CONCAT	
INITCAP	
LENGTH	
LOWER	
LPAD	
LTRIM	
REPLACE	
RPAD	
RTRIM	

PostgreSQL Character Function	Notes
SPLIT	
SPLIT_PART	
STRPOS	
SUBSTR	
SUBSTRING	
TRANSLATE	
TRIM	
UPPER	

PostgreSQL Conversion Function Support

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

PostgreSQL Conversion Function	Notes
CAST	PostgreSQL does not preserve trailing spaces when casting CHARs to VARCHARs, so results may differ when a federated data source or TDV is set to honor trailing spaces. PostgreSQL data types have the following restrictions for the CAST function: • Maximum CHAR length is 2000. • Maximum VARCHAR length is 4000. • Maximum numeric precision (p) is 38. • Maximum numeric scale (s) is 38.

PostgreSQL Conversion Function	Notes
FORMAT_DATE	
PARSE_TIMESTAMP	
TO_CHAR	
TO_DATE	
TO_NUMBER	
TO_TIMESTAMP	

PostgreSQL Date Function Support

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

PostgreSQL Date Function	Notes
CLOCK_TIMESTAMP	
CURRENT_DATE	
CURRENT_TIMESTAMP	
DATE_TRUNC	
DAY	
MONTH	
YEAR	Not supported.
DATE_PART	This is supported in Postgres version 10 and later.

PostgreSQL Date Function	Notes
EXTRACT(DOW FROM)	Push supported. The tdv function returns value in range [1, 7]. But when the function is pushed to Postgresql, the returned value is in range [0, 6].

PostgreSQL Numeric Function Support

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

PostgreSQL Numeric Function	Notes
ABS	
ACOS	
ASIN	
ATAN	
ATAN2	
CEIL	
CEILING	
COS	
COT	
DEGREES	
EXP	
EXTRACT	From a date.

PostgreSQL Numeric Function	Notes
FLOOR	
LOG	
MOD	
PI	
POW	Not supported: mixing string, number, or NULL with BIT, NCHAR, or NVARCHAR.
POWER	
RADIANS	
RANDOM	
ROUND	
SIGN	
SIN	
SQRT	
TAN	
TRUNC	

PostgreSQL Specific Properties

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

Port	Port number for the data source to connect with the host. The Port number for PostgreSQL is 5432.
Enable Bulk Export/Load	Takes advantage of PostgreSQL COPY command.
Enable PostgreSql dblink	Check to enable database links to improve performance if you plan to use this data source for data caching or data ship optimization. If you check this box, add one or more database links by specifying the database link name and path of the data source for each link

PostgreSQL Time Function Support

TDV supports the time functions listed in the table below for PostgreSQL.

PostgreSQL Time Function	Notes
EXTRACT	From TIMESTAMP or INTERVAL_DAY.
NOW	
TIMEOFDAY	

References

Refer to the following Guides for further details about the capabilities of the data source:

Capabilities	Section
Query Engine	User Guide, Chapter ***TDV Query Engine Optimizations***
Data ship	User Guide, Chapter ***Data Ship***

Capabilities	Section
	Performance Optimization
Caching	User Guide, Chapter ***TDV Caching***
Performance Optimization	User Guide, Chapter ***Performance Tuning***
TDV Massively Parallel Processing Engine	User Guide, Chapter ***Configuring the TDV MPP Engine***
Kerberos	Administration Guide Chapter ***Configuring Kerberos***

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 orservice 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 theTIBCO 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.

TIBCO, TIBCO logo, Two-Second Advantage, TIBCO Spotfire, TIBCO ActiveSpaces, TIBCO Spotfire Developer, TIBCO EMS, TIBCO Spotfire Automation Services, TIBCO Enterprise Runtime for R, TIBCO Spotfire Server, TIBCO Spotfire Web Player, TIBCO Spotfire Statistics Services, S-PLUS, and TIBCO Spotfire 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.