TIBCO StreamBase® Release Notes
Software Release 10.3.0

Important information

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 DOCUMENTATION 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 DOCUMENTATION 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 documentation contains confidential information that is subject to U.S. and international copyright laws and treaties. No part of this documentation may be reproduced in any form without the written authorization of TIBCO Software Inc.

TIBCO, TIBCO StreamBase, TIBCO StreamBase Studio, TIBCO StreamBase Server, TIBCO StreamBase EventFlow, TIBCO Live Datamart, TIBCO LiveView Desktop, TIBCO LiveView Web, TIBCO LiveView Web Standard Edition, TIBCO LiveView Web Enterprise Edition, TIBCO Rendezvous, TIBCO Enterprise Message Service, TIBCO ActiveSpaces, and TIBCO FTL are trademarks or registered trademarks of TIBCO Software Inc.

Enterprise Java Beans (EJB), Java Platform Enterprise Edition (Java EE), Java 2 Platform Enterprise Edition (J2EE), and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle Corporation in the U.S. and other countries.

All other product and company names and marks mentioned in this documentation 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 DOCUMENTATION 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 DOCUMENTATION 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 DOCUMENTATION. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENTATION AT ANY TIME.

THE CONTENTS OF THIS DOCUMENTATION 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.

StreamBase Release Notes

TIBCO StreamBase 10.3.0

Date: 24 Jul 2018

These are the release notes for TIBCO StreamBase® 10.3.0 with internal tracking number 261916.

1. Before You Install

Before you install TIBCO StreamBase® and TIBCO® Live Datamart software:

1. Read the New and Noteworthy section of the Release Notes for important information on migrating from previous StreamBase releases.

2. Confirm that your target system matches the hardware and software requirements discussed on Supported Configurations.

3. Read the Installation Guide.

2. Documentation Starting Points

Getting Started provides a tutorial introduction to TIBCO StreamBase® and TIBCO® Live Datamart.

Concepts Overview describes the concepts and terminology of TIBCO StreamBase® and TIBCO® Live Datamart.

3. New and Noteworthy in StreamBase 10.3.0

This section describes the new and updated features of TIBCO StreamBase® 10.3.0.

3.1. New Features

TIBCO StreamBase® 10.3.0 adds the following updates and new features:

3.1.1. New Unified Security Model

Security Model Changes

The security model in StreamBase 10.3.0 is a significant upgrade over previous releases, providing greater security without sacrificing ease of use during development. See StreamBase Security Model for details. New features include:

- Unified configuration settings for StreamBase and Live Datamart communication, and for StreamBase Runtime node access communication.

- New support for an OpenID Connect (OIDC) authentication realm. OIDC is an identity layer on top of the OAuth 2.0 protocol which allows single sign-on for clients to verify your identity based on the authentication performed by an authorization server. The following identity providers are currently supported:
  - Google Identity Platform
  - Auth0
  - Microsoft Azure Active Directory

- New support for a Kerberos Authentication Realm.

- Separation of configuration for secure server and secure client communication.

Security Behavior Changes

Some behavior changes are the result of this upgrade:

- In previous releases, when communicating with a node on another machine (including on virtual machines and in Docker containers), you could use the `guest` user as a default entry point. Starting with release 10.3.0, at minimum you must configure either a list of trusted hosts or a password for the default user in order to communicate with the remote machine. This is described in Security Configuration Defaults.

- The **epadmin install node** command's `trustedhosts` parameter was removed. Instead, you must define trusted hosts using the TrustedHosts root object in a Security configuration file.

3.1.2. New Maven Artifact Management

New Repository Handling

StreamBase 10.3.0 introduces a new way to manage StreamBase-provided JAR files and other artifacts for consumption by its Maven build system.

- In StreamBase 10.2, artifacts were installed in the `sdk/maven/repo` directory of your StreamBase installation directory. Maven POM files generated by StreamBase Studio in 10.2 were configured to read from this directory as one of its Maven repositories.

- In StreamBase 10.3.0 and later, artifacts are installed in an archive file. StreamBase Studio then merges the contents of the archive into your local Maven repository, which is usually the `.m2/repository` folder of your home directory. On every start, Studio checks the status of the local repository and merges any missing files back in. This allows the POM files generated by Studio 10.3 to use only standard Maven conventions.

Do Not Open StreamBase 10.2 Workspaces with Studio 10.3

TIBCO has always recommended against trying to use a Studio workspace from a previous release in a new major release. However, this recommendation is especially important for the StreamBase 10.2 to 10.3 transition because of 10.3's new repository handling described above.

DO NOT OPEN StreamBase 10.2 project folders in Studio from StreamBase 10.3. See Migrating from 10.2 to 10.3 for further information.

New epdev Command

StreamBase 10.3.0 introduces the **epdev** command, which allows you to develop on a machine without a Studio installation, as in a devops environment. Use this command to unpack the shipped Maven artifacts into a machine's local repository. You can also use **epdev** to correct repository errors, to write the StreamBase artifacts to a non-default repository, to configure a development laptop for offline use, and other tasks. See the epdev reference page or **epdev --help** for further information.

TIBCO_EP_HOME Variable No Longer Used

As part of the artifact handling change described above, StreamBase 10.3 no longer requires or uses the TIBCO_EP_HOME environment variable, either by Studio or command-line tools. There is no harm in leaving the variable set, but it is not used.

However, StreamBase 10.1 and 10.2 still requires that variable. If you are maintaining old and new Studio versions, make sure any shell startup files or Windows scripts set that variable correctly for use only with StreamBase installations before 10.3.

3.1.3. New Studio Support for StreamBase Applications

Running Applications in Studio

StreamBase 10.3.0 introduces support for running StreamBase Applications in StreamBase Studio. Previous StreamBase 10 releases required Apps to be run from the command line with the **epadmin** command. While **epadmin** deployment remains the gold standard for production deployment, launching Apps in Studio allows you to test or debug an application consisting of two or more fragments.

While a StreamBase Application is running, you can stop and restart it from the right-click menu in the Debug view.

New Run Configuration Type

This release introduces a new Run Configuration type named StreamBase Application. The minimum requirement for this configuration is an application archive (whose filename usually ends with `*ep-application.zip`). This can be generated in Studio with the wizard described next, or on the command line with **mvn package**.

Once you have an application archive file, right-click the filename, and select Run As → Run Configurations. In the Run Configurations dialog, create a new StreamBase Application configuration and use the first Browse button on the Main tab to locate the your `*ep-application.zip`. Then click Run. See Run StreamBase Application in Studio.

New Wizard to Generate Application Archives

As an aid for running StreamBase Apps either with **epadmin** or in Studio, this release includes a new wizard that first generates fragment archives for all dependent fragments, and then generates an archive for the StreamBase App.

In previous releases, you had to run Maven Install on the fragment projects you wanted to encapsulate in a StreamBase Application project, and then Maven Install again on the App project. The new wizard reduces those steps to a single menu selection.
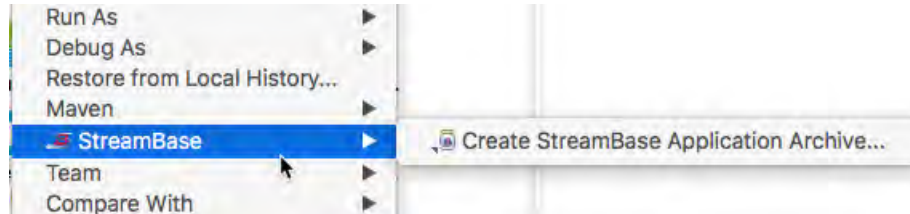

Workspace Resolution for Dependencies

When attempting to locate projects to be designated as Maven dependencies, Studio now looks in the current Studio workspace as well as the local Maven repository. This allows Studio to locate projects without having to run Maven Install on them first.

3.1.4. New Unit Test Support for Fragments

Previous StreamBase 10 releases supported running Java unit tests that load a designated EventFlow fragment, send tuples to it, and look for expected results. Previously, these tests could be run with Run As → Maven Test or on the command line with **mvn test**. Those methods are still available, but this release adds a new options to the Run As menu.
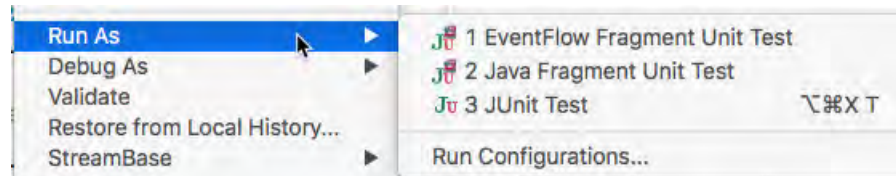
The two Fragment Unit Test options run the selected Java unit test and connect the results to Studio's JUnit view to see the results as green success or red failure bars. You can select the project's `src/test/java` folder to run all unit test files underneath, or select a package name, if your project has more than one, to run all tests with that package name.

3.1.5. New epadmin Syntax and Features

epadmin New Syntax

The **epadmin** command now supports Gnu-style parameter syntax with two hyphens prefixing parameter names. You can enter as few letters of each parameter's name that preserves that parameter's uniqueness among the current **epadmin** target's supported parameters.

Syntax from previous releases is still supported. In this format, parameters are entered in full without abbreviations and without a prefixed double hyphen. You cannot mix old and new syntax in the same command. See the **epadmin** reference page for details.

epadmin help Expanded

The usage text for **epadmin** commands is greatly expanded and clarified, and you no longer need to connect to a node before seeing usage text. Start with **epadmin help**, **epadmin help targets**, and **epadmin help globals**. Thorough usage text is available for all targets by using **epadmin help** *targetname*.

epadmin Command Error Message Enhanced

The returned error message was enhanced for cases when you enter a partial, non-unique parameter for a given **epadmin** command. The error message now returns the name of conflicting parameters, and also suggests a partial name for the parameter that could enable a successful command execution.

epadmin verify services Command Added

This command verifies that the UDP broadcast mechanism used by StreamBase Runtime discovery services is working normally on the local subnet. See **epadmin --help verify services** for details.

epadmin tail logging Command Added

This command emits new messages as they are added to an engine's log stream. For further notes, see Tail Logging.

## 3.1.6. New Studio Usability Enhancements

New HOCON Configuration Chooser

This release shows a new HOCON type selection dialog when you run File → New → StreamBase HOCON Configuration File. This dialog organizes HOCON file types in root object categories and lets you enter a search string to narrow the list to matching root objects.
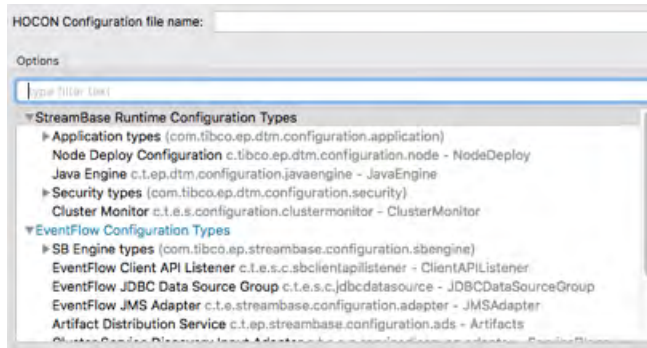
New Console View Output Level Settings

The Console view was enhanced to allow you to set the level of information displayed when running fragments or applications. When the Console view has fragment or application output to display, a new Format menu in the Console view toolbar sets one of three options: Verbose, Concise, or Minimal.
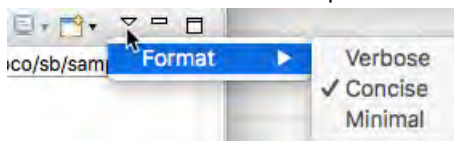
The Format menu is a shortcut for selecting one of the same preferences from the Launching panel of Studio Preferences. Concise is Studio's default in StreamBase 10.3.0; in previous releases, the Console always showed Verbose format. For a description of these settings, see Launching Panel.

New Clusters View Display Options

In previous 10.x releases, the Properties panel of the Clusters view was always placed on the right side of the view. Starting with this release, Studio provides an option to locate the Properties panel on the right or below the Clusters, Nodes, and Engines panel. The new default is below.
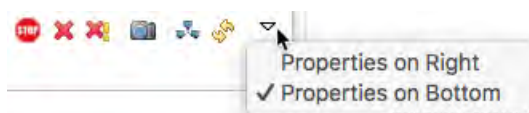
Fragment Dependencies Now Shown in Project Explorer

When one fragment project has a Maven dependency on other fragment projects, a new folder named StreamBase Fragment Project Dependencies appears in the Project Explorer view to show the dependency relationships.

New JDBC Data Source Connection Options

This release includes the following JDBC configuration connection options:
• Connection pooling allows all Query operators that use the same data source to share the same connection.

- Multiple JDBC connections can now be configured in a single HOCON configuration file by using the new JDBC-DataSourceGroup root object.

See StreamBase JDBC DataSource Configuration and JDBC Data Source Overview for a discussion of using the JDBC Table data construct and one or more Query operators to access an external JDBC data source from an EventFlow module.

Copying Schema to Clipboard Now Supported

Studio now provides the ability to copy the schema of a selected input stream, named schema, or arc to the system clipboard, in either JSON or XML format. This feature allows you to copy a schema ready to paste into the Create Artifact dialog of the Artifact Management Server. See Exporting Schemas to Clipboard for details.

3.1.7. New Expression Language Function

New Function to Calculate Great Circle Distance

This release adds the great_circle_distance() function to the StreamBase expression language.

3.1.8. New Platform Support

Additions to Operating System Support

StreamBase and Live Datamart are now supported for build or staging servers and for production deployment on:
- Oracle Linux 7
- Microsoft Windows Server 2016

Oracle JDK Version Updated

The Oracle JDK version bundled with StreamBase and Live Datamart is now JDK 1.8.0_u171.

Command Line **mvn** Command Shipped with StreamBase

StreamBase 10.3.0 and later ships with the open source **mvn** command installed. If you use a StreamBase Command Prompt on Windows or a StreamBase-configured shell on macOS or Linux, the **mvn** command is in the PATH, as is the **java** command required to run **mvn**.

Windows Service Now Supported

StreamBase now supports running a StreamBase application as a Windows service, including support for using event logging. See Configuring Windows as a Service for details.

Docker Support Enhanced

OpenShift Container Orchestration Environments: Cluster auto-scaling is supported using Kubernetes Horizontal Pod Autoscale (HPA) 1.7. See Supported Docker Environments.

Jetty Upgraded

In StreamBase Studio, the Jetty web server version was upgraded to 9.3.23.v20180228. The upgrade applies to, for example, the TIBCO StreamBase® Web Server Request adapter and the TIBCO StreamBase® Web Server Response adapter, as well as for developing your own applications that require Jetty.

3.1.9. New Adapters

New CME iLink Order Routing Adapter

The adapter named CME Order Routing in previous releases was removed and replaced with the CME iLink Order Routing adapter. This new adapter has a different usage model, which requires applications using the previous version to be modified accordingly. The new adapter includes:
- Support for CME's Secure Logon procedure
- Support for Convenience Gateway
- Support for Market Segment Gateway (with failover support)

See CME iLink Order Routing Adapter for configuration details and CME Order Routing Adapter Sample for a sample of adapter usage.

New CME Market Data Protocol Adapter

The adapter named CME Market Data Platform FIX/FAST Input adapter in previous releases was removed and replaced with the CME MDP3 adapter. The new adapter allows StreamBase to receive data from a CME Market Data Protocol (MDP3)

connection, and now supports explicit IP address binding for use on multi-homed systems, without requiring additional operating system configuration.

See CME MDP3 Adapter for configuration details and CME MDP3 Adapter Sample for a sample of adapter usage.

eFTL Adapters Now Available

This release includes the following new eFTL adapters:

- TIBCO StreamBase® Subscribe Adapter for TIBCO eFTL®, which allows a StreamBase application to receive eFTL messages. See TIBCO eFTL Subscribe Input Adapter for details.

- TIBCO StreamBase® Publish Adapter for TIBCO eFTL®, which allows a StreamBase application to publish eFTL messages. See TIBCO eFTL Publish Output Adapter for details.

New Kafka Adapter Added

The TIBCO StreamBase® Output Adapter for Apache Kafka Consumer Commit allows StreamBase applications to perform commit operations for a given consumer operator. The commit can be a commit of all topics and partitions or be a more granular approach and specify the topic, partition, and offset to commit. See Kafka Consumer Adapter for more information.

New Web Server Adapters Added

This release includes the following new web adapters:

- TIBCO StreamBase® Web Input Adapter allows the StreamBase application to receive JSON formatted tuples via an HTTP POST request. For adapter details, see Web Input Adapter.

- TIBCO StreamBase® Web Output Adapter allows the system to stream tuple data in JSON format to one or more connected web sockets. For adapter details, see Web Output Adapter.

See Web Streams Adapter Sample for a sample of the adapters' usage.

New Task Scheduler Adapter

The TIBCO StreamBase® Task Scheduler adapter allows Studio to fire events at specific intervals. See Task Scheduler Adapter Samples for sample usage and Task Scheduler Adapter to learn more about its properties.

3.2. Changes in Functionality

The following changes in functionality occurred between releases 10.2.1 and 10.3.0.

3.2.1. Studio Changes

Select Configuration Files from Other Projects

In previous releases, in a Run Configuration dialog for an EventFlow fragment, you could select from among the configuration files in the `src/main/configurations` and `src/test/configurations` folders for the current project only. Starting with 10.3.0, you can select configuration files from other projects in the workspace as well. This feature is a convenience for short-term testing only. Configuration files designated this way do not automatically become part of the current project's archive.

Properties View, General Tab

In the Properties view for all global Java operators and adapters, in the General tab, previous releases combined the operator or adapter line with the operator or adapter class name into a single long line. Starting with 10.3.0, the fully qualified class name is on a line of its own. As before, you can select and copy the class name for pasting elsewhere, such as into a Java JUnit file.

3.2.2. Installation Changes

Some Windows Command-line Tools were Moved

Only on Windows installations, certain command-line tools were moved to the `bin64` directory of your StreamBase installation. If the PATH contains both the `bin` and `bin64` directories (as provided by the StreamBase Command Prompt), you can expect no changes. The commands that moved are:

jsbadmin   sb-config
jsbc   sbfeedsim

jsbclientgen   sbprofile
sbcipher

3.2.3. HOCON Configuration Changes

Creating Fragment Projects No Longer Generates fragment.conf

In previous releases, creating a fragment project in Studio would automatically generate a `fragment.conf` file in `src/main/configurations` as a configuration placeholder. Studio no longer generates this file for new fragment projects. In most cases such `fragment.conf` files contained nothing meaningful to the project being created. You can still generate a Fragment Definition HOCON file manually and configure it as required.

Trusted Hosts Configuration Supports New Formats

StreamBase and Live Datamart trusted host configurations now support partially qualified domain names and IPv6 addresses. See Trusted Hosts for more information.

Fragment Specifications Now Require Only the Top-Level Fragment

In previous releases, if you used the `fragments` object under a `nodeTypes` definition in the ApplicationDefinition root object of the application HOCON type, you had to separately list the top-level fragment and all of its dependent fragments (if any). Starting with StreamBase 10.3.0, you only need to specify the top-level fragment in a node type's fragment list, and transitive dependencies are automatically included.

memorySizeBytes Property Now Has a Minimum Value

The `memorySizeBytes` property of the ApplicationDefinition configuration type configures the size of transactional memory for a node. That property now has a minimum value of 512 MB, which has also been its default value. Previous releases allowed you to configure a lower value for this property.

Sparse Partitions are now Foreign Partitions

Sparse partitions were renamed to *foreign partitions*. Other than terminology, however, the feature is unchanged. The Sparse Audit Output column in the **epadmin display partition** command was changed to Foreign Partition Audit Option.

3.2.4. epadmin Changes

export logging Now Requires **enginename**

The **epadmin export logging** command's `enginename` parameter is now required (it was previously optional). Release 10.3.0 removes support for exporting logging configuration for all engines using a single command.

display configuration No Longer Shows **Principal**

The Principal column was removed from the **epadmin display configuration** command output. The node audit log can now show the user who made changes to the configuration state.

In the **realm** Target, the --realm Parameter is Renamed

The **--realm** parameter is now **--name** throughout the **epadmin** *verb* **realm** command.

3.2.5. Operator Updates

Model Operators Include New Input Schema Option

The input schema for the JPMML, H2O Model Evaluator and Spark/MLlib Model Evaluator operators now optionally allows you specify the model name field against which the input tuple should execute. If the field is missing or null, then existing functionality of executing against all loaded models occurs.

Distributed Router Operator Change

Error handling changed in the Distributed Router operator when using the Broadcast routing policy. Previous releases would return an error if a target node reported an error. StreamBase now ignores these errors. Change your Distributer Router configuration accordingly.

3.2.6. Adapter Updates

CME Adapters Replaced

Both adapters that provide connections to Chicago Mercantile Exchange resources were replaced in this release. See New Adapters above.

LogLogic Adapter Name Changed

The TIBCO StreamBase® Adapter for TIBCO® LogLogic Unity's name changed to TIBCO StreamBase® Adapter for TIBCO® LogLogic. The adapter's functionality remains unchanged. See TIBCO LogLogic Query Input Adapter for adapter information.

HDFS Adapter Suite Now Supports Amazon S3 Connectivity

The HDFS adapter suite was enhanced to support connectivity to an Amazon S3 file system. The enhanced adapters are:
- TIBCO StreamBase® Binary File Reader For Apache Hadoop Distributed File System (HDFS)
- TIBCO StreamBase® File System Adapter for HDFS
- TIBCO StreamBase® CSV File Reader For Apache HDFS
- TIBCO StreamBase® File Reader For Apache Hadoop Distributed File System (HDFS)
- TIBCO StreamBase® Regular Expression File Reader For Apache Hadoop Distributed File System (HDFS)
- TIBCO StreamBase® Binary File Writer For Apache HDFS
- TIBCO StreamBase® CSV File Writer For Apache Hadoop Distributed File System (HDFS)
- TIBCO StreamBase® File Writer Adapter for Apache HDFS
- TIBCO StreamBase® XML File Writer for Apache Hadoop Distributed File System (HDFS)

New samples illustrate how to access the Amazon S3 file system using the TIBCO StreamBase® File Writer for Apache Hadoop Distributed File System (HDFS) and TIBCO StreamBase® File Reader for Apache Hadoop Distributed File System (HDFS). See Amazon S3 File System Access via HDFS Adapters.

Spotfire Automation Services Adapter API Changed

To provide support for three versions of the TIBCO Spotfire automations services, the TIBCO StreamBase® Output Adapter for TIBCO Spotfire® Automation Services now features an API version drop-down list and deprecates the Previous Version Compatibility check box.

With this change the default API version is now 1.0, which may break existing implementations. Use version 0.1 to return to functionality that works with Spotfire 7.4 and below, or use version 0.2 to return to functionality that works with Spotfire releases 7.5 to 7.12. For more information regarding this adapter's properties, see Spotfire Automation Services Output Adapter.

FTL Subscribing Input Adapter Improved for Security and Latency

The TIBCO StreamBase® Input Adapter for TIBCO FTL® was enhanced to include support for secure connections, durables and inline mode. Additionally, you can now programmatically change the ContentMatcher string for the subscription. For more information about these and other properties, see TIBCO FTL Subscribing Input Adapter.

UDP Send and Receive Adapters Enhanced

The TIBCO StreamBase® Input and Output adapters for UDP now include a new property to specify which network interface adapter to bind. If left blank (default), the adapters function as before and bind to all interfaces.

Additionally, the adapters now both include a Data Transformer tab which allows you to specify multiple options to set up how the received or sent data will be transformed.

See UDP Sender Output Adapter and UDP Receiver Input Adapter for details regarding the Sender Output and Receive Input adapters, respectively.

Multiple Kafka Consumer and Producer Adapter Changes

Most of the individual options located in the Advanced tab of the Properties view for the Kafka Consumer and Producer Adapters were removed in favor of a grid where you can enter key-value pairs for advanced options. These options are based on the adapters being upgraded to Kafka version 11:

https://kafka.apache.org/0110/javadoc/org/apache/kafka/clients/producer/ProducerConfig.html
https://kafka.apache.org/0110/javadoc/org/apache/kafka/clients/consumer/ConsumerConfig.html
https://kafka.apache.org/0110/javadoc/constant-values.html

If an advanced feature of a Kafka adapter was set with StreamBase 7.7.2 or earlier, it may no longer be a Kafka option, and thus now missing from the Properties view. If an advanced feature does still exist, you must now use the configuration key and value to reset the option.

Consumer groups are now supported as part of the upgrade of the Kafka Consumer adapter Kafka version 11. As a result of the upgrade, simple consumer functionality is now deprecated.

The Kafka Consumer adapter's advanced configuration settings are no longer restricted to be set per adapter. The adapter now supports advanced configuration options per subscription tuple.

See Kafka Producer Adapter and Kafka Consumer Adapter for further information about Kafka adapter functionality.

Tuple2Avro and Avro2Tuple Operators Can Now Specify Avro Schema

The Tuple2Avro and Avro2Tuple operators can now specify an Avro schema when converting between binary Apache Avro records and tuples. If no value is specified, the operator continues to generate an Avro schema from the input/output tuple schema as before.

IBM WebSphere MQ Adapter Suspend and Resume Behavior Changed

When suspended, the TIBCO StreamBase® Input Adapter for IBM MQ Series now stops emitting tuples. The message thread continues to run and receive messages but those messages are discarded. For adapter details, see IBM WebSphere MQ Input Adapter.

Web Reader Adapter URL Encoding Options Added

The TIBCO StreamBase® Web Reader adapter previously included a single option to specify URL encoding. Two options are now available:
• Decode HTML results — decodes strings before output.
• URL Encode Post Data — specifies URL-encoding the PostData value.

See Web Reader Input Adapter for adapter option details.

Web Server Adapters Enhanced

This release includes additional functionality for the context path of the TIBCO StreamBase® Web Server Request adapter. Now, the context path can also contain path variables in the form of {*varname*}. Additionally, the Web Server Request Adapter can now convert:
• JSON requests to tuples
• XML requests to tuples
• Binary requests to a series of tuples containing the request POST data

The Web Server Response Adapter can now convert:
• Tuples to JSON responses
• Tuples to XML responses
• Tuples to POST responses

See Web Server Request Adapter and Web Server Response Adapter for details.

File Reader Adapter Now Reads Hierarchical Files

The TIBCO StreamBase® File Reader adapter now allows you to specify a folder or filename for reading files. If you specify a folder, all files contained therein that match the given glob pattern are read. For more information regarding this adapter's properties, see File Reader Input Adapter.

JMS Consumer Operator Adds Synchronous Tuple Emission Option

The JMS Consumer operator now includes a new setting: Emit Tuples Synchronously. This optional setting causes the operator to wait until the tuple is fully processed before dequeueing the next message from the JMS server. In this way, implicit backpressure can be applied to the upstream server if the incoming traffic becomes unmanageable.

3.3. Deprecated and Removed Features

StreamBase 10.3.0 has the following deprecations.

3.3.1. RHEL 6 Deprecated

Support for Red Hat Enterprise Linux 6.x is deprecated and is expected to be removed in a future release.

3.3.2. Deprecated HOCON Configuration Features

The following HOCON configuration features are deprecated but are still supported for backward compatibility:

Local Admin Authentication Realm

In the Security HOCON type, the entire Local Admin Authentication Realm root object is deprecated. Use Local Authentication Realm instead.

principals Property

The `initialPrincipals` property replaces the deprecated `principals` property when configuring a Local Authentication Realm root object of the Security HOCON type.

JDBCDataSource Root Object

The `JDBCDataSource` root object in the JDBC Data Source configuration type is deprecated and replaced with the `JDBC-DataSourceGroup` root object, which allows you to configure more than one JDBC data source definition in the same file.

3.3.3. Subclipse Deprecated

Studio compare features no longer support Subclipse, which is now removed. Instead, use Subversive.

3.3.4. sbcipher Command Deprecated

Password and other string data enciphered with the **sbcipher** command is still deciphered and used in StreamBase 10.3.0 and later. Going forward, however, use the **epadmin encrypt secret** and master secret features of the security model introduced in 10.3.0 and described in Encrypting Sensitive Configuration Data.

3.3.5. XML Normalizer Operator Deprecated

The XML Normalizer operator is now deprecated. Applications that use the operator continue to compile and run, but the operator no longer appears in the Operators and Adapters dialog. Documentation for the operator and its sample were removed.

3.3.6. Deprecated Visual Studio 2010 for C++ Removed

Visual Studio 2010 for C++, which was announced as deprecated in the StreamBase 10.2.1 New and Noteworthy, is now removed from the Windows SDK. TIBCO recommends using Visual Studio 2012 as the baseline for C++ client development.

3.3.7. sbproxy Command Removed

Previous StreamBase releases relied on the **sbproxy** command to provide a secure listener for StreamBase server. Starting with release 10.3.0, StreamBase server supports a native TLS listener. Thus, the **sbproxy** command was removed as inoperable with release 10.3.0's new security model.

3.4. Migration and Compatibility

Migration Guide Provided

This release provides a guide for customers migrating from previous StreamBase 10.x releases in Migrating from 10.2 to 10.3.

JMS and TIBCO EMS Adapter Configuration Files New Format

StreamBase 10.3.0 introduces a new HOCON configuration type for configuring JMS and TIBCO EMS server connections, using the JMSAdapter root object in the com.tibco.ep.streambase.configuration.adapter type.

Prior to Release 10.3.0, configuration for JMS and EMS operators was stored in `src/main/resources/adapter-configura-tions.xml`. For migrating projects to 10.3.0, the `adapter-configurations.xml` file still works, with no modifications necessary. Edits made in 10.3.0 with Studio's JMS Editor automatically converts the JMS configuration to a new file in HOCON format, and then comments out the same configuration from the XML file. See JMS and TIBCO EMS Configuration Editor.

StreamBase Studio Now Based on Eclipse 4.7.3

StreamBase Studio, StreamBase Manager, and LiveView Desktop are now based on Eclipse 4.7.3. In Studio, you may notice that some menu items are in different places. Some formerly generic dialog buttons labelled OK and Close now state the action to be performed instead.

LiveView Desktop users may be prompted to upgrade their workspaces because the Eclipse 4.7 workspace format has changed. StreamBase Studio silently updates the workspace format. However, TIBCO strongly recommends importing your projects from previous releases into a new, unused Studio 10.3 workspace rather than trying to reuse a previous release's workspace.

Commented Substitution Variables Inside Quoted Strings No Longer Ignored

Prior to StreamBase 10.3.0, substitution variables commented AND inside a quoted string were ignored in configuration files of HOCON type. As of 10.3.0, they are no longer ignored.

4. Limitations Resolved in 10.3.0

This section provides a list of errors corrected in release 10.3.0 of TIBCO StreamBase®.

Fixed in 10.3.0	
Number	**Resolution**
TIBCO StreamBase® 10.3.0 incorporates all fixes resolved in the release 10.0 series through 10.3.0 and all fixes resolved in the release 7.7 series through 7.7.4.	
SB-44502	If a Query Table or Table Schema was based on a named schema, then the partitioning drop-down menu in the table editing toolbar might have been disabled. It is now enabled, thereby allowing you to specify a partition when appropriate.
SB-44475	In previous releases, the `querytable` administration target would only work for query tables in a top-level module. This was fixed and query tables in nested modules are now correctly handled.
SB-44457	Specifying one or more substitution name-value pair definitions in a fragment launch configuration (as opposed to a naming a file of substitution name-value pairs), would not actually use those definitions during the launch. This was fixed.
SB-44202	In previous releases, having an EventFlow fragment that contained a Query Table with no data distribution policy prevented the running a JUnit test. Now unit tests do run in these circumstances, with an INFO message issued. Tuples are added to the Query Table, but the data is not partitioned.
SB-44168	An issue prevented the OSI PI adapter from loading the .NET OSI PI DLL. This was fixed for the entire OSI PI adapter suite.
SB-43899, SB-29388	When logging off the StreamBase HP FIX engine, the FIX adapter was put into a state that prevented the logon command from working when attempting to log on again. This was fixed.
SB-43848	An issue prevented the Kafka Producer adapter from closing on error, regardless of the option for doing being set. This was fixed.
SB-43797	The Kafka Consumer and Producer adapters were fixed to address an issue that prevented suspend and resume behavior from working correctly.
SB-43648, SB-43458	Artifact deployment could fail when the PMML model operator pulled in artifacts from an Artifact Management Server (AMS) via the Artifact Distribution Service (ADS) — and when the default container was explicitly set for deployment in AMS. Now, the default container name is no longer required for deployments.
SB-43783	This release fixes an issue that prevented new models from completely replacing existing ones, when loaded from the JPMML Model Evaluator Operator's control port.
SB-43621	There was a problem where duplicate entries could mistakenly appear in the fragments drop-down list in the Manual Input view's Output Streams and Input Streams views. This problem could occur immediately after the launch dialog for an EventFlow fragment launch disappeared, when running a second EventFlow fragment launch. This was fixed.
SB-43566	The `Operator.evaluate` Java API would previously return a value of zero for data types int, long, double and Boolean if their values were in fact null. This was fixed so that null is returned in these cases.
SB-43479	An issue caused the JPMML Model Evaluator operator to issue NaN errors when running PMML models containing special characters and spaces in the data field names. The operator now unwraps any field name that uses escaped identifier syntax before trying to compare to any field names in the actual models. This ensures the field names will now match. The unwrapped name now also appears in the log for easier reference.
SB-43464	This release fixes an issue that prevented EventFlow applications from launching if an undefined data distribution policy was defined for a query table in a unit test case. A warning is now generated. Tuples are still added to the query table in shared memory, but data is not partitioned.
SB-43413	All field types were updated on the TERR operator to check for a data index range. Now, a null value is returned when trying to access an index beyond what is in the TERR object.
SB-43293	This release fixes an issue in Studio's module drag and drop capability that prevented file names from converting to fully qualified names, when dropped into nested projects.
SB-43161	Previously, MQTT Input adapters using the same shared connection would output messages for all topics — whether they subscribed to the topics or not. Now, each adapter outputs messages only for the topics to which they are subscribed.
SB-43111	Previously, Studio would return a ResourceNotFoundException typecheck error if an EventFlow module containing a Decision Table application called for another module's Decision Table file. This was fixed.

Fixed in 10.3.0	
Number	**Resolution**
SB-43032	This release fixed a URI parameter encoding issue in the Web Reader Input adapter.
SB-42764	The WITS adapter typecheck could fail if the `stringNullValue` adapter configuration option was missing. Now, the adapter no longer emits an NPE error if the WITS adapter configuration is missing that option.
SB-42746	The Web Reader adapter's status port schema was only set during typecheck and thus was required to be set at initialization, with the runtime version of the schema. As a result, the adapter was not passing capture fields to the error output port. The Web Reader adapter's status port now works correctly with capture fields.
SB-42394	In previous releases, a collision between virtual network interfaces and the StreamBase Runtime caused (a) extremely slow node installation times on Windows and/or (b) when more than one fragment was run at the same time, nodes would self-destruct, leaving no evidence in log files. Virtual network interfaces are installed as a part of virtualization hypervisor software such as Oracle VirtualBox, Microsoft Hyper-V, or VMware Workstation. This collision occurred only on Windows, and was triggered by installing Docker for Windows, which requires and enables Hyper-V. The cause of the collision was identified and corrected, and no longer occurs with StreamBase releases 10.3.0 and later. See Programs with Incompatibilities for workarounds if you are running an earlier StreamBase 10 release.
SB-42250	Model artifacts deployed from the Artifact Management Server to a StreamBase cluster via the Artifact Distribution Service (ADS) would occasionally fail to deploy on all nodes. The underlying deadlock was fixed to address this issue.
SB-41775	In cases when the Kafka driver detected a TCP/IP connection drop, the Kafka Consumer adapter would emit an inaccurate warning message describing the cause. The Kafka Consumer now correctly reports the failure error when the adapter exceeds the maximum number of retries.

5. Known Limitations

This section describes known limitations in the current release of TIBCO StreamBase®. Each item includes a tracking number, description, and workarounds when possible.

Known limitations for using the EventFlow debugger are described separately, in the Limitations, Tips, and Tricks section of the Using the EventFlow Debugger page.

SB-44282	Error when building Docker images in Studio on Mac.	
	Description	An error results when following the steps on a Mac to build a Docker image containing a StreamBase Application. The Console view references a failure to locate the **docker-credential-osxkeychain** command.
	Workaround	To make sure `/usr/local/bin` is in Studio's PATH, start Studio by typing `sbstudio &` at a StreamBase-configured Terminal shell prompt.
SB-44167	iPhone connected to Mac can cause node launch failures.	
	Description	Conditions: an iPhone is connected to a Mac with a USB-to-Lightning cable while Personal Hotspot is NOT enabled on the iPhone. Under these conditions, StreamBase Runtime nodes can fail to start.
	Workaround	Unplug the iPhone, or enable Personal Hotspot (whether or not you actually use Personal Hotspot).
SB-42478	Cluster Monitor fails to start in the same cluster as monitored.	
	Description	Attempting to start the Cluster Monitor in the same cluster as other applications can fail.
	Workaround	Run the Cluster Monitor feature in a separate cluster from the one containing the applications being monitored.
SB-42410	Duplicate engine log files.	
	Description	Duplicate engine log files exist in both the node directory's top level and in its `logs` subdirectory. The other user visible log files are now all located in the `logs` subdirectory.
	Workaround	None. The engine log files in the node directory's top level can be ignored. The same information is available in the engine log files found in the `logs` subdirectory. The duplicate log files are expected to be removed in a future release.

SB-42324, SB-30773, SB-30846	Inconsistent containerName behavior.	
	Description	If you define a StreamBase container name with the `containerName` parameter of the `sbengine` HOCON configuration type, the runtime behavior is different for fragments run in Studio compared to **epadmin** launches. Studio places the EventFlow module in a container named *default*, following standards set in StreamBase 7.x, whereas **epadmin** honors the configuration-set container name.
	Workaround	None.
SB-42408	ADS: load failure if artifact is not replicated on each node.	
	Description	Using dynamic partitioning for the Artifact Distribution Service (ADS) will cause duplicate key errors when loading artifacts, if the artifacts are not replicated to all nodes in a cluster.
	Workaround	Ensure that artifacts are replicated to all nodes in a cluster.
SB-42407	Artifacts reported as missing and duplicate created after cluster shape change.	
	Description	The Artifact Distribution Service (ADS) can report existing artifacts as missing and create duplicate artifacts after a node is added or removed from a cluster, if high availability is not configured.
	Workaround	Configure high availability for ADS if the cluster shape is going to be changed.
SB-42093	Reloading the entire node configuration when attempting to configure the cluster monitor fails.	
	Description	Attempting to reload a cluster monitor configuration embedded in a node deploy configuration file using the **epadmin load configuration** command quietly does nothing.
	Workaround	Create a separate configuration file with just the cluster monitor configuration and load and activate that configuration.
SB-41896	Studio Support Wizard screen capture can fail on Mac.	
	Description	The StreamBase Support Wizard screenshot feature preview may not display properly under certain multi-monitor and/or Retina display hardware conditions.
	Workaround	Use the Support Wizard to guide you to continue; the Support zip file will still contain the screenshot that you can verify in the archive prior to sending to TIBCO. For more information on the Support Wizard, see StreamBase Support Wizard for Troubleshooting.
SB-41099	Calling Help via management API throws null pointer exception	
	Description	Executing `help` commands using the `DtmCommand.execute` method in the management API fails with an exception.
	Workaround	None. This issue is expected to be resolved in a future release.
SB-35326	Shutting down at debugger breakpoint causes engine to fail with an error.	
	Description	Stopping a JVM with a breakpoint set in the debugger can cause the JVM to exit with an error that requires the node to be removed and reinstalled.
	Workaround	Ensure that no breakpoints are set in a debugger when shutting down a JVM.
SB-32239	Bogus `FATAL Invalid response` trace seen when stopping a node.	
	Description	When deploying fragments, a misleading fatal error message may be returned when terminating the fragment or the node on which it is deployed. For example: `MESSAGE: dtm : FATAL(RQ:A.dstest): Invalid response: <ready/>`
	Workaround	This message can be safely ignored.
SB-30948	Trouble using the TERR operators on macOS.	
	Description	On macOS, EventFlow modules that use the TERR operators with TERR versions before 4.4 fail to start even when using the Use Embedded TERR check box. This includes the TERR operator samples shipped with StreamBase.
	Workaround	If you can, update to TERR 4.4 or later, which is included with StreamBase 10.2.1 and later. If you must use a TERR version before 4.4, there are extra configuration steps on macOS, compared to Windows and Linux. These steps are detailed in the Operating System Configuration for TERR sections of the pages that describe the TERR Operator and the TERR Predict Operator.
SB-30898	On Windows only, virtual network adapters slow fragment startup with StreamBase before 10.3.0.	

	Description	If you have Oracle VM VirtualBox® or Microsoft Hyper-V installed, and one or more of your virtual machines has its virtual network adapter activated, this slows down the launch of EventFlow and LiveView fragments into a StreamBase Runtime node by several minutes.
	Workaround	Upgrade to StreamBase 10.3.0 or later, where this issue is corrected. If you must use an earlier release, temporarily disable the network adapter feature of your VirtualBox VMs as described in Windows and Virtual Network Interfaces.
SB-30270	Failure to start Studio or Manager by double-clicking on macOS Sierra or High Sierra	
	Description	MacOS releases 10.12 Sierra and 10.13 High Sierra have a security model that is stricter than previous releases. As part of their security, these releases automatically block un-signed, downloaded application bundles with `.app` extension from being able to run. Un-fortunately, this affects the application bundles for StreamBase Studio 10.x and Stream-Base Manager 10.x.
	Workaround	See the list of suggested workarounds in Starting StreamBase Studio on macOS Sierra.
SB-29050	Error when debug stepping into or stepping over a Decision Table operator.	
	Description	When using EventFlow debugging, stepping into or over a Decision Table operator is not supported.
	Workaround	Set your breakpoints upstream or downstream. Also, consider using the features of the Decision Table Analyzer to get a debug-like view of how the operator responds to different input tuples.
SB-28856	Deleting decision table rows does not re-number rule IDs.	
	Description	When using the Decision Table editor to delete and then add a new decision table rule, the rule ID does not renumber. For example, by deleting row 5 and adding a new row, the rule ID numbering is now 1...4, 6.
	Workaround	None.
SB-28669	Error message: Grows beyond 64KB.	
	Description	In rare cases, compiling small StreamBase applications can produce an error message containing the words `Grows beyond 64KB`.
	Workaround	Contact TIBCO StreamBase Technical Support for a workaround for this issue.
SB-28598	Multi-container StreamBase applications cannot use the same ActiveSpaces metaspace.	
	Description	StreamBase applications deployed with multiple containers receive `metaspace already exists` errors when attempting to connect to the same TIBCO ActiveSpaces metaspace from different containers.
	Workaround	None.
SB-28452	Delta stream arc limitations in the EventFlow Debugger.	
	Description	If a delta stream arc is a direct connection between a Query Table and an Output Stream, and if that arc has a Debugger breakpoint, the breakpoint is not hit at during debugging with the EventFlow Debugger. Similarly, while debugging, stepping over to try to reach that arc does not suspend on that arc.
	Workaround	Insert a no-op Map operator between the Query Table and the Output Stream. Then a breakpoint on the arc between the Query Table and the Map does work, and stepping over reaches that arc as well.
SB-26629	RPM installation method issues on RHEL 7	
	Description	StreamBase and Live Datamart are supported on Red Hat Enterprise Linux version 7.x. The archive installer is the recommended installation method. If you use the RPM installation method, RHEL 7 issues several warning messages during installation that complain about file removal failures, and issues an overall warning that RHEL 7 "may not be compatible with this version of StreamBase." Despite these warnings, the installation does succeed. The only consequence is that `rpm -e` fails as an uninstallation method.
	Workaround	Remove StreamBase directories manually to uninstall an RPM installation on RHEL 7. Oth-erwise, all StreamBase and Live Datamart processes run as expected on RHEL 7.
SB-26927	**Error message** `NamedSchema is not a function that can be called`.	

	Description	In compiling the updating expression for a dynamic variable where the expression references a named schema, if a field in the tuple presented by the updating stream also has the same name as the named schema used in the expression, the field is instead evaluated as a function. This raises the error *NamedSchema is not a function that can be called*. This compilation behavior and language ambiguity was introduced along with the function data type.
	Workaround	The required migration step is to rename either the field or the named schema to eliminate the ambiguity.
SB-26237	Constant definitions used as component properties are not used at runtime.	
	Description	StreamBase Studio properly typechecks defined constants when used in operator properties, but at runtime substitutes their names instead of their values.
	Workaround	Do not use defined constants when specifying operator properties.
SB-26017	Asynchronous timeouts for Gather operators in parallel applications can be inaccurate	
	Description	High availability applications that include Gather operators may incorrectly compute timeout intervals because the operator relies on now() instead of calling out to the parallel scheduling service.
	Workaround	None
SB-22297	Commented-out error input streams consume error tuples.	
	Description	An error input stream, despite being disabled with the EventFlow comment-out feature, still catches errors thrown, but consumes any incoming errors and fails to pass them on.
	Workaround	To temporarily disable error streams, copy them to a placeholder module and remove them from the active module.
SB-21843	In the TableAccessor API, ClassCastException is thrown if not using a Tuple from TableAccessor#createRow.	
	Description	If methods in the `com.streambase.sb.operator.TableAccessor` interface that accept a Tuple are not provided a Tuple retrieved from `TableAccessor#createRow`, a ClassCastException may be thrown by the API.
	Workaround	Always use `TableAccessor#createRow` when using TableAccessor API methods.
SB-20449	Display issues with years before 2 AD.	
	Description	The timestamp data type can hold timestamp values in the range $-(2^{62})$ to $(2^{62})-1$, and can successfully display any future date in the next 146 million years and any past date to around year 2 AD. However, the year value for dates earlier than 2 AD may not display correctly.
	Workaround	None.
SB-18294	Union operator: sub-field order of tuple fields must match with strict options.	
	Description	When using the Union operator with its Strict union or Declared schema options, the order of sub-fields of tuple fields must match exactly. This is in contrast with the default Loose union option, where strict sub-field ordering is no longer required.
	Workaround	None.
SB-16594	Extract as Module feature copies items instead of importing them.	
	Description	The Extract as Module feature makes a copy of the source module's definitions and named schemas, and places the copies in the newly created module. Best practice would be to import those features from the source module to the new module, and to copy all imports from the source to new modules.
	Workaround	After using Extract as Module, manually convert the copied definitions and named schemas to imports.
SB-14133	On Linux, **sb-config --cxx** returns a setting for G++ newer than 4.5.	
	Description	On Linux, in StreamBase samples that build C++ code, some of the example Makefiles show **sb-config --cxx** as the correct usage to determine the G++ compiler to use. On newer Linux distributions, this command returns the distribution's default G++ compiler, which may not be supported for building StreamBase client code, and **sb-config** does not honor the CC and CXX environment variables.
	Workaround	On such systems, install GCC and G++ 4.5, and edit the Makefiles to specify `CXX=g++-4.5`.

SB-13398	Filter operators with very large predicate settings may result in a runtime EvalException error.	
	Description	Filter operators with several output ports and complex expressions to evaluate for each port may result in a `com.streambase.sb.runtime.exceptions.EvalException` error when attempting to run the application that contains such a filter.
	Workaround	Add another Filter operator and split the predicate settings between the two Filter operators.
SB-13113	Application output view does not display UTF-8 characters.	
	Description	Both StreamBase Server and StreamBase Studio are fully capable of processing and displaying Unicode character sets, but neither is set to do so by default.
	Workaround	Follow the instructions in Unicode Support to configure StreamBase Server, Studio, and any StreamBase client applications to process the UTF-8 character set.
SB-10792	The Query operator converts a JDBC data type differently when using explicit schemas than when using database data types.	
	Description	When explicit schemas are used, the Query operator does not use database metadata to determine the mapping between JDBC data types and StreamBase data types. The JDBC data type `Types.TIME` is normally converted to an interval timestamp since it has a time, but no date component. But when explicit schemas are used, there is no database type information, so `Types.TIME` is converted to an absolute timestamp.
	Workaround	None.
SB-10486	Hierarchical data restrictions.	
	Description	Patterns by value do not support using the tuple data type for their on-value portion. Similarly, tuple subfields cannot be used as the index or primary key of a Query Table.
	Workaround	Design schemas so that the intended sort field or key is a top-level numeric or timestamp field, as described in StreamBase Pattern Matching Language.
SB-10363	No StreamBase data type corresponding to Sybase timestamp.	
	Description	When used with Sybase data sources, both EvenfFlow and StreamSQL fail to convert timestamps, resulting in typecheck errors. This is because the Sybase `timestamp` is based on the Sybase `varbinary` data type.
	Workaround	In EventFlow or StreamSQL code that will be used for Sybase data sources, use the Sybase `datetime` data type instead of `timestamp`.
SB-6862	Not all databases have native blob support.	
	Description	For example, Microsoft SQL Server and Sybase do not have blob support. For these databases, you can write blobs into varbinary and possibly other types. But when blobs are read from a select statement, they are returned as strings, not blobs.
	Workaround	Add a Map operator and convert the string to a blob.
SB-5882	Tuple dimensions do not reset counter when another dimension closes a window early.	
	Description	A tuple-based dimension that closes due to other dimensions (for example, after another dimension times out) does not reset its counter. When this happens, the next window will close based on the total number of tuples received, resulting in a partially full dimension.
	Workaround	None.
SB-4559	Join timeout is not asynchronous.	
	Description	Join timeout in the Join operator is not asynchronous. That is, the Join operator does not remove a tuple from its buffer until the timeout has passed *and* another tuple arrives on *the same* input stream. Tuples received on one input stream do not affect the buffer of the other input stream.
	Workaround	None.

6. Previous Release Notes

See Previous Release Notes for the history of new features, updates, and resolved issued in previous TIBCO StreamBase® 10 releases.