TIBCO Flogo® Connector for Apache Kafka

User Guide

Version 1.3.3 | January 2025

Contents

Overview

Apache Kafka is a distributed messaging system, providing fast, highly scalable, and redundant messaging through a publisher-subscriber model. By using TIBCO Flogo® Connector for Apache Kafka connectivity, you can design the flows to send and receive the records.

Flogo® Connector for Apache Kafka connectivity supports OpenTelemetry. For more information, see the *"Tracing Apps by Using OpenTelemetry Collector"* topic of TIBCO Flogo® EnterpriseTIBCO Flogo® app documentation.

Flogo Connector for Apache Kafka supports Flow Limit. For more information, see the *Environment Variables* topic of TIBCO Flogo® EnterpriseTIBCO Flogo® app documentation.

For more information about Apache Kafka, see Kafka documentation.

Configuring a Kafka Client Connection

To publish or subscribe messages, you must first configure a Apache Kafka client connection. The Apache Kafka client connection contains the parameters required to connect to the Apache Kafka cluster. The Apache Kafka client connection is used by all the activities in the Apache Kafka category.

Note: As TIBCO Cloud Integration supports ports in the range 1024-9999, ensure that the open port of the local broker configuration or the hosted Apache Kafka service is in this valid range.
Example of a valid service URL used in connection configuration for brokers fields: `kafka-3507f830- johndoe-eb63.abccloud.com:2819`

Procedure

1. On the global navigation bar, click the **Connections** and perform one of the following actions:
    - To add a connection for the first time, click the **Create connection** link and

then select the **Apache Kafka Client Configuration** card. You can search for a connector card by entering the connector name in the search field.

- To create subsequent connections, click the **Create** button.

2. In the Apache Kafka Client Configuration dialog, enter the connection details. For field descriptions, see Kafka Connection Details.
3. Click **Save Connection**.

Kafka Client Configuration Details

To establish the connection successfully, you must have configured the Apache Kafka instance.

The **Apache Kafka Client Configuration** dialog contains the following fields.

Condition Applicable	Field	Description
N/A	Connection Name	The unique name for the connection you are creating. This name is displayed in the **Connection** dropdown list for all the TIBCO Flogo® Connector for Apache Kafka activities.
N/A	Description	A short description of the connection
N/A	Brokers	A comma-separated list of host and port pairs (host:port) for establishing the initial connection with the Kafka cluster.
N/A	Auth Mode	Select one of the following authentication types to connect with the Kafka cluster: • **None:** To establish the connection without authentication • **SASL/PLAIN:** To use Simple Authentication Security Layer (SASL) PLAIN authentication

Condition Applicable	Field	Description
		• **SSL:** To use Secure Socket Layer (SSL) authentication • **SASL/SCRAM-SHA-256:** To connect to the Kafka cluster configured for SASL/SCRAM with hash SHA-256 functions • **SASL/SCRAM-SHA-512:** To connect to the Kafka cluster configured for SASL/SCRAM with hash SHA-512 functions • **SASL/OAUTHBEARER:** To use Simple Authentication Security Layer (SASL) with OAuth Bearer functions
Applicable only when **SASL/PLAIN**, **SASL/SCRAM-SHA-256**, , **SASL/SCRAM-SHA-512** is selected in the **Auth Mode** field.	User Name Password	Enter the credentials (required for authentication).
Applicable only when **SASL/PLAIN** is selected in the **Auth Mode** field.	Security Protocol	Select one of the following security protocols: • **SASL_PLAINTEXT:** Nonsecure connection to the broker • **SASL_SSL:** Default. Secure connection to the broker
Applicable only when **SASL/PLAIN**, **SSL, SASL/SCRAM-SHA-256, SASL/SCRAM-SHA-512or SASL/OAUTHBEARER:** is selected in the **Auth Mode** field.	Client Certificate	A Privacy Enhanced Mail (PEM) encoded client certificate file for mutual authentication.
	Client Key	A PEM encoded private key file for mutual authentication.

Condition Applicable	Field	Description
	CA or Server Certificate	A PEM encoded private key file for server authentication.
Applicable only when **SASL/OAUTHBEARER** is selected in the **Auth Mode** field.	Client ID	The public identifier for an application used to support the OAuth client credentials grant type.
	Client Secret	The secret, known only to an application and an authorization server, associated with the clientID and used to support the OAuth client credentials grant.
	Scope	The scope to reference in the call to the OAuth server.
	Token URL	The URL for the OAuth 2.0 issuer token endpoint.
N/A	Connection Timeout	The amount of time in seconds to wait for the initial connection. **Default:** 30 seconds
N/A	Retry Backoff	The amount of time in milliseconds to wait for the leader election to occur before retrying. **Default:** 250 milliseconds
N/A	Max Retry	The number of attempts to retry metadata requests when the cluster is in the middle of a leader election. **Default:** 3 attempts
N/A	Refresh Frequency	The amount of time in seconds after which metadata is refreshed. **Default:** 40 seconds

Condition Applicable	Field	Description
N/A	Use Schema Registry	Enables you to use the Avro schema with a Schema Registry by selecting **True**. If you select **True**, you must specify a **Schema Registry URL**. You can use the Confluent Schema Registry or the TIBCO Schema Registry. Depending on which Schema Registry you want to use, the **Schema Registry URL** needs to be changed. Basic authentication is used to connect to the registry if a **Username** and **Password** are provided. **Note:** • If **True** is selected, the schema displayed in the **Schema for Avro Value** field in the following tabs is read-only. If **False** is selected, you can create or modify the schema on both these tabs. ◦ Input Settings tab of Kafka Producer ◦ Output Settings tab of Kafka Consumer Trigger • SSL authentication is not supported. **Default value:** False
Applicable only when **True** is selected in the **Use Schema Registry** field.	Schema Registry URL	The URL used to connect to a Schema Registry. For the Confluent Schema Registry, the URL must be in the following format: `http://<host>:<port>` For the TIBCO Schema Registry, the URL

Condition Applicable	Field	Description
		must be in the following format: `http://<host>:<port>/schema/v1` **Note:** For TIBCO Schema Registry, use the FTL realm URL.
	Username	(Optional) The username to access the Schema Registry with basic authentication.
	Password	(Optional) The password to access the Schema Registry with basic authentication.

Note: Client Certificate, Client Key & Server Certificate supports app property with the following formats:

- encrypted certificates
- base64 encoded value
- certificates as a files

Connecting to TIBCO Cloud Messaging - Kafka

Configuring Kafka Client Configuration through tcm-config.yaml

You can configure Kafka Client Configuration by using a `tcm-config.yaml` file as follows:

- **Broker:** Kafka broker url from the `tcm-config.yaml` file
- **Auth Mode:** SASL/PLAIN
- **USERNAME:** Kafka username from the `tcm-config.yaml` file
- **PASSWORD:** token: <tcm_authentication_key> (TCM authentication key must be prefixed by 'token:')
- **SECURITY_PROTOCOL**: SASL_SSL

For more information about downloading the `tcm-config.yaml` file, see the "TIBCO Cloud Messaging Client Configuration File" section in TIBCO Cloud™ Messaging documentation.

Support For Azure EventHub

For Azure Event Hub, below connection fields are required:

- Brokers - <azure event hub namespace>.servicebus.windows.net:9093
- Authentication Type - SASL/PLAIN
- Username - $ConnectionString
- Password - <Connection String -Primary Key>
- Security Protocol - SASL_SSL

For more information about how to configure Event Hub, click the link:

https://learn.microsoft.com/en-gb/azure/event-hubs/event-hubs-create

Kafka Consumer Trigger

Consumer Trigger receives records from a specified topic in the Apache Kafka cluster.

Trigger Settings

On the **Settings** tab, define the Apache Kafka connection and its details as given in the following table:

Condition Applicable	Field	Description
N/A	Apache Kafka Client Configuration	Apache Kafka client configuration to be used.
N/A	Topic	The topic where Apache Kafka cluster stores a stream of

Condition Applicable	Field	Description
		records. **Note:** If you are using multiple topics, separate them using commas.
	Handler Settings	
N/A	Assign Custom Partition	Enable to assign custom partition. When enabled, the subscriber can subscribe to only a single topic(First value in case of comma separated topics) mentioned in the topic and single partition Id of that topic. To subscribe on combination of topic and partition, you need to add multiple handlers. Default is false, to keep Kafka's default consumer group model. When set to true, it allows you to add Partition Id from where Kafka consumer is able to read messages from particular Partition Id.
N/A	Partiton ID	Enter the Partition ID to consume the messages.
N/A	Consumer Group ID	The group ID for the consumer group.
N/A	Value Deserializer	Select the type of record value to be received from the dropdown list: • **String** • **JSON** • **Avro**
Applicable only when **Avro** is selected in	Subject	A list of all registered subjects in your Schema Registry. A subject refers to the name under which a schema is registered. Select the subject to be used.

Condition Applicable	Field	Description
the **Value Serializer** field.		
Applicable only when **Avro** is selected in the **Value Serializer** field.	Version	Version of the subject (registered name of schema) registered. Select the version of the subject to be used.
N/A	Commit Interval	The time interval in which a consumer offset commits to Apache Kafka. **Default:** 5,000 milliseconds
N/A	Initial Offset	Select one of the following options: • **Newest:** To start receiving published records since the consumer is started • **Oldest:** To start receiving records since the last commit • **Seek By Offset:** This option is available when Assign Custom Partition is set to True. In Apache Kafka, seeking by offset refers to the process of moving a consumer's position to a specific offset within a topic partition. This is useful when you want to read messages from a particular point in a partition, whether for reprocessing or starting from a known offset. • **Offset:** Seek messages starting from specified integer offset. • **Seek By Timestamp:** This option is available when Assign Custom Partition is set to True. Seeking by

Condition Applicable	Field	Description
		timestamp in Kafka allows you to start consuming messages from a specific point in time. • **Timestamp:** Seek messages starting from specified RFC3339 timestamp.
N/A	Fetch Min Bytes	Minimum size of data that the server sends on fetch request.
N/A	Fetch Max Wait	The maximum amount of time that the server would block before answering a fetch request if there is not sufficient data to satisfy the requirement immediately that you have configured in the **Fetch Min Bytes** field.
N/A	Heartbeat Interval	Time in milliseconds to send heartbeats to consumers. Heartbeats are used to ensure that the consumer's session remains active and to facilitate rebalancing when consumers join or leave a group. **Note:** Heartbeat interval must not be more than one-third of the session time.
N/A	Session Timeout	The consumer sends periodic heartbeats to the server indicating its liveness to the broker. If no heartbeats are received by a broker before the session times out, the broker removes this consumer from the group and initiates a rebalance.

Output Settings

Condition Applicable	Field	Description
N/A	Headers	Record headers to be received. Only String datatype value is supported.

Condition Applicable	Field	Description
		Note: Headers are supported in Apache Kafka version 0.11.0 and later.
Applicable only when **JSON** is selected in the **Value Serializer** field on the **Trigger Settings** tab.	Schema for JSON value	The JSON schema for the Apache Kafka record value.
Applicable only when **Avro** is selected in the **Value Serializer** field on the **Trigger Settings** tab.	Schema for Avro Value	The Avro schema for the Apache Kafka record value. Depending on the **Subject** and **Version** selected, the schema is displayed here. **Note:** This field is read-only if **Use Schema Registry** in the **Apache Kafka Client Configuration** dialog is set to **True**. Otherwise, you can provide the schema using this editor.

Output

Condition Applicable	Field	Description
Applicable only when **JSON** is selected in the **Value Serializer** field.	jsonValue	Data structure based on the JSON schema that you have configured in the **Output Settings** section.
Applicable only when **Avro** is selected in the **Value Serializer** field.	avroData	Data structure based on the Avro schema that you have configured in the **Output Settings** section.
N/A	partition	Partition number of the record.
N/A	offset	Offset of the record.
N/A	topic	Name of the topic.

Condition Applicable	Field	Description
N/A	key	Key value.
Applicable only when **String** is selected in the **Value Deserializer** field on the **Settings** tab.	stringValue	String value to be received.
N/A	headers	Header value to be received.

Kafka Producer

Apache Kafka producer activity sends a record to a specified topic or channel in the Kafka cluster.

Settings

On the **Settings** tab, you can define the Apache Kafka connection and its details as given in the following table:

Condition Applicable	Field	Description
N/A	Apache Kafka Connection	Select the connection that you want to use from the dropdown list.
N/A	Acks Mode	Select one of the following acknowledgment modes from the dropdown list: • **None:** To receive no acknowledgment on record delivery • **Leader:** To receive an acknowledgment on record delivery from the leader • **All:** To receive acknowledgment on record delivery from leaders and all in-sync replicas

Condition Applicable	Field	Description
Applicable only when **All** is selected in the **Ack Mode** field.	Ack Timeout	The amount of waiting time in milliseconds to receive confirmation.
N/A	Partitioner	Select one of the following partitioners from the dropdown list: • **Hash:** This is the default mode. When you select the **Hash** mode and the **key** field is not set on the **Input** tab, then a random partition is selected, otherwise, a key hash is used for calculating the partition to ensure all records with the same key are sent to the same partition. • **Manual:** When you select the **Manual** mode, the value configured in the **partition** field of the **Input** tab is used to select the partition. **Note:** If the **Manual** mode is selected, then the **partition** field is mandatory on the **Input** tab.
N/A	Compression Type	Select a compression type: **None**, **GZIP**, or **LZ4**.
N/A	Value Serializer	Select the type of record value to be sent: • **String** • **JSON** • **Avro**
Applicable only when **Avro** is selected in the	Subject	A list of all registered subjects in your Schema Registry. A subject refers to the name under which a schema is registered. Select the subject to be used.

Condition Applicable	Field	Description
Value Serializer field.		
Applicable only when **Avro** is selected in the **Value Serializer** field.	Version	Version of the subject (registered name of schema) registered. Select the version of the subject to be used.
N/A	Max Request Size	The maximum size of buffered records that can be sent in one request. **Default value:** 1048576 bytes
N/A	Max Messages	The maximum number of records that can be sent in a single broker request.
N/A	Frequency	The frequency of sending buffered records in milliseconds. **Default value:** 1000

Input Settings

Condition Applicable	Field	Description
N/A	Headers	Header record to be sent. Only the String datatype value is supported. **Note:** Headers are supported in the Apache Kafka version 0.11.0 and later.
Applicable only when	Schema for JSON	The JSON schema for the Apache Kafka record value.

Condition Applicable	Field	Description
JSON is selected in the **Value Serializer** field on the Settings tab.	value	
Applicable only when **Avro** is selected in the **Value Serializer** field on the Settings tab.	Schema for Avro Value	The Avro schema for the Apache Kafka record value. Depending on the **Subject** and **Version** selected, the schema is displayed here. **Note:** This field is read-only if **Use Schema Registry** in the **Apache Kafka Client Configuration** dialog is set to **True**. Otherwise, you can provide the schema using this editor.

Input

Condition Applicable	Field	Description
N/A	topic	Name of the topic.
N/A	partition	Existing partition number where the record is to be sent.
N/A	key	Optional key value.
Applicable only when **String** is selected in the **Value Serializer** field.	stringValue	String value to be sent.
Applicable only when **JSON** is selected in the **Value Serializer** field.	jsonValue	Data structure based on the JSON schema that you have configured on the **Input Settings** tab.
Applicable only when **Avro** is selected in the **Value Serializer** field.	avroData	Data structure based on the Avro schema that you have configured on the **Input Settings** tab.
N/A	headers	Header value to be sent.

Output

Condition Applicable	Field	Description
N/A	topic	Name of the topic.
N/A	partition	Partition number of the record to send.
N/A	offset	Offset of the record.

Loop

For information on the **Loop** tab, see the Using the Loop Feature in an Activity section in the TIBCO Flogo® EnterpriseTIBCO Flogo® app documentation.

Retry on Error

Using the Retry on Error tab, you can set the number of times the flow tries to run the activity on encountering an error that can be fixed on retrial. The errors such as waiting for a server to start, intermittent connection failures, or connection timeout can be fixed on retrial.

You can set the count and the interval in one of the following ways:

- Manually type the value in the mapper.
- Map the value from the previous Activity.
- Select a function from the list of functions.
- Map app property to override the values.

Field	Description
Count	The number of times the flow should attempt to run the activity. This value must be an integer.
Interval (in millisecond)	The time to wait in between each attempt to run the activity. This value must be an integer.

Note: The Count and Interval fields are mandatory. By default, the values are set to 0.

Kafka Offset Commit

Apache Kafka Offset Commit activity notifies Kafka Consumer Trigger to commit given an offset. This is useful in the case that you want offsets to be committed as soon as the record is processed in the flow. By default, offsets are committed only when the flow is successfully completed.

Note: This activity can be used only with Kafka Consumer Trigger.

Input

Condition Applicable	Field	Description
N/A	topic	Name of the topic.
N/A	partition	Existing partition number where the record is to be sent.
N/A	offset	Offset of the record.

Loop

For information on the **Loop** tab, see the Using the Loop Feature in an Activity section in the TIBCO Flogo® Enterprise TIBCO Flogo® app documentation.

TIBCO Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join 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 documentation for this product is available on the TIBCO Flogo® Connector for Apache Kafka Product Documentation page.

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.

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, go to TIBCO Community.

Legal and Third-Party Notices

SOME CLOUD SOFTWARE GROUP, INC. (“CLOUD SG”) SOFTWARE AND CLOUD SERVICES EMBED, BUNDLE, OR OTHERWISE INCLUDE OTHER SOFTWARE, INCLUDING OTHER CLOUD SG SOFTWARE (COLLECTIVELY, “INCLUDED SOFTWARE”). USE OF INCLUDED SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED CLOUD SG SOFTWARE AND/OR CLOUD SERVICES. THE INCLUDED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER CLOUD SG SOFTWARE AND/OR CLOUD SERVICES OR FOR ANY OTHER PURPOSE.

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

TIBCO, the TIBCO logo, the TIBCO O logo, and Flogo are either registered trademarks or trademarks of Cloud Software Group, Inc. in the United States and/or other countries.

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. You acknowledge that all rights to these third party marks are the exclusive property of their respective owners. Please refer to Cloud SG’s Third Party Trademark Notices (https://www.cloud.com/legal) for more information.

This document includes fonts that are licensed under the SIL Open Font License, Version 1.1, which is available at: https://scripts.sil.org/OFL

Copyright (c) Paul D. Hunt, with Reserved Font Name Source Sans Pro and Source Code Pro.

Cloud SG 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 a specific version of Cloud SG software on a specific operating system platform.