TIBCO Flogo® Connector for Apache Kafka User's Guide

Software Release 1.1.1
February 2019

Two-Second Advantage®

TIBCO®

Important Information

Contents

TIBCO Documentation and Support Services

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. To access the latest documentation, visit https://docs.tibco.com.

Documentation for TIBCO Flogo® Connector for Apache Kafka is available on the TIBCO Flogo® Connector for Apache Kafka Product Documentation page.

Product-Specific Documentation

The following documents for this product can be found on the TIBCO Documentation site:

- *TIBCO Flogo® Connector for Apache Kafka Installation*
- *TIBCO Flogo® Connector for Apache Kafka User's Guide*
- *TIBCO Flogo® Connector for Apache Kafka Release Notes*

How to Contact TIBCO Support

You can contact TIBCO Support in the following ways:

- For an overview of TIBCO Support, visit http://www.tibco.com/services/support.
- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the TIBCO Support portal at https://support.tibco.com.
- 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 https://support.tibco.com. 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 https://community.tibco.com.

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, you can design the flows to send and receive the records.

For information about how to use Apache Kafka, see Kafka documentation.

Configuring a Kafka Client Connection

To use TIBCO Flogo® Connector for Apache Kafka, 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.

Prerequisites

Before you create a client connection, familiarize yourself with Apache Kafka. For details about how to use the Apache Kafka product, see the Kafka Documentation.

Procedure

1. On the TIBCO Flogo® Enterprise page, click the **Connections** tab and perform one of the following actions:
   - To add a client connection for the first time, click the **Apache Kafka Client Configuration** card. You can search for a connector card by typing the connector name in the search field.
   - If you have existing connections and want to add a new connection, click the **Add Connection** link.
2. In the Apache Kafka Client Configuration dialog box, enter the connection details. For field descriptions, see the Kafka Connection Details topic.
3. Click **Save**.

Kafka Client Configuration Details

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

The **Apache Kafka Client Configuration** dialog box 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** drop-down list for all the TIBCO Flogo® Connector for Apache Kafka activities.
N/A	Description	A short description of connection
N/A	Brokers	A comma-separated list of host and port pair (host:port) for establishing the initial connection with the Kafka cluster.

Condition Applicable	Field	Description
Applicable only when **SASL/PLAIN** is selected in the **Auth Mode** field.	Auth Mode	Select one for following authentication type to establish the connection with Kafka cluster: • **None:** To establish the connection without authentication • **SASL/PLAIN:** To use Simple Authentication Security Layer (SASL) PLAIN authentication • **SSL:** To use Secure Socket Layer (SSL) authentication
	User Name	The user name for authentication.
	Password	The password for authentication.
Applicable only when **SASL/PLAIN** or **SSL** 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.
	CA or Server Certificate	A PEM encoded private key file for server authentication.
	Connection Timeout	The amount of time in seconds to wait for the initial connection. **Default value:** 30 seconds
N/A	Retry Backoff	The amount of time in milliseconds to wait for leader election to occur before retrying. **Default value:** 250 milliseconds
N/A	Max Retry	The number of attempts to retry metadata request when the cluster is in the middle of a leader election. **Default value:** 3 attempts
N/A	Refresh Frequency	The amount of time in seconds after which metadata is refreshed. **Default value:** 40 seconds

Kafka Consumer Trigger

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

Configuration

On the **Configuration** 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 Client Configuration	Kafka client configuration to be used.
N/A	Topic	The topic where Kafka cluster stores streams of record.
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 drop-down list: **String** or **JSON**
N/A	Commit Interval	The time interval in which a consumer offset commits to Kafka. **Default Value:** 5000 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
N/A	Fetch Min Bytes	Minimum size of data that 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 immediately satisfy the requirement that you have configured in the **Fetch Min Bytes** field.
N/A	Heartbeat Interval	Time in milliseconds to send heartbeats to consumer. Heartbeats are used to ensure that the consumer's session remains active and to facilitate rebalancing when consumers join or leave a group.  Heartbeat interval must not be more than one-third of the session time.

Condition Applicable	Field	Description
N/A	Session Timeout	The consumer sends periodic heartbeats to server indicating about 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.  Headers are supported in the Apache Kafka version 0.11.0 and later.
Applicable only when **JSON** is selected in the **Value Serializer** field on the **Configuration** tab.	Schema for JSON value	The JSON schema for the Kafka record value

Output

Condition Applicable	Field	Description
Applicable only when JSON is selected in the **Value Serializer** field.	jsonValue	Complex data structure based on JSON 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
N/A	key	Key value
Applicable only when String is selected in the **Value Deserializer** field on the **Configuration** 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.

Configuration

On the **Configuration** 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 you want to use from the drop-down list.
N/A	Acks Mode	Select one of the following acknowledgement modes from the drop-down list: • **None:** To receive no acknowledgement on record delivery • **Leader:** To receive an acknowledgement on record delivery from the leader • **All:** To receive acknowledgement on record delivery from leaders and all in-sync replicas
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	Compression Type	Select a compression type: **None**, **GZIP**, or **LZ4**.
N/A	Value Serializer	Select the type of record value to be sent: **String** or **JSON**.
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 `String` datatype value is supported.  Headers are supported in the Apache Kafka version 0.11.0 and later.
Applicable only when **JSON** is selected in the **Value Serializer** field.	Schema for JSON value	The JSON schema for the Kafka record value.

Input

Condition Applicable	Field	Description
N/A	topic	Name of the topic.
N/A	partition	Partition number of the record to send.
N/A	key	Optional key value.
Applicable only when **String** is selected in the **Value Serializer** field.	stringValue	String value to be send
Applicable only when **JSON** is selected in the **Value Serializer** field.	jsonValue	Complex data structure based on JSON 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.

Iterator

Use the **Iterator** tab to iterate a certain piece of logic multiple times. If you leave this tab blank, the activity is executed only once. For more information about Iterator, see "Using the Iterator in an Activity" in the TIBCO Flogo® apps documentation.

Kafka Offset Commit

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

This activity can be used only in conjuction with Kafka Consumer Trigger.

Iterator

Use the **Iterator** tab to iterate a certain piece of logic multiple times. If you leave this tab blank, the activity is executed only once. For more information about Iterator, see "Using the Iterator in an Activity" in the TIBCO Flogo® apps documentation.