



Red Hat AMQ Streams 2.1

Using the AMQ Streams Kafka Bridge

Use the AMQ Streams Kafka Bridge to connect with a Kafka cluster

Red Hat AMQ Streams 2.1 Using the AMQ Streams Kafka Bridge

Use the AMQ Streams Kafka Bridge to connect with a Kafka cluster

Legal Notice

Copyright © 2022 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

The AMQ Streams Kafka Bridge provides a RESTful interface for HTTP-based clients to interact with a Kafka cluster.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	8
CHAPTER 1. KAFKA BRIDGE OVERVIEW	9
1.1. RUNNING THE KAFKA BRIDGE	9
1.2. KAFKA BRIDGE INTERFACE	10
1.2.1. HTTP requests	10
1.3. SECURING CONNECTIVITY TO THE KAFKA CLUSTER	11
1.4. SECURING THE KAFKA BRIDGE HTTP INTERFACE	11
1.5. REQUESTS TO THE KAFKA BRIDGE	11
1.5.1. Content Type headers	11
1.5.2. Embedded data format	12
1.5.3. Message format	12
1.5.4. Accept headers	13
1.6. CORS	13
1.6.1. Simple request	14
1.6.2. Preflighted request	14
1.7. CONFIGURING LOGGERS FOR THE KAFKA BRIDGE	15
CHAPTER 2. KAFKA BRIDGE QUICKSTART	17
2.1. DOWNLOADING A KAFKA BRIDGE ARCHIVE	17
2.2. CONFIGURING KAFKA BRIDGE PROPERTIES	17
2.3. INSTALLING THE KAFKA BRIDGE	18
2.4. PRODUCING MESSAGES TO TOPICS AND PARTITIONS	19
2.5. CREATING A KAFKA BRIDGE CONSUMER	25
2.6. SUBSCRIBING A KAFKA BRIDGE CONSUMER TO TOPICS	26
2.7. RETRIEVING THE LATEST MESSAGES FROM A KAFKA BRIDGE CONSUMER	27
2.8. COMMITTING OFFSETS TO THE LOG	28
2.9. SEEKING TO OFFSETS FOR A PARTITION	28
2.10. DELETING A KAFKA BRIDGE CONSUMER	29
CHAPTER 3. AMQ STREAMS KAFKA BRIDGE API REFERENCE	31
3.1. OVERVIEW	31
3.1.1. Version information	31
3.1.2. Tags	31
3.1.3. Consumes	31
3.1.4. Produces	31
3.2. DEFINITIONS	31
3.2.1. AssignedTopicPartitions	31
3.2.2. BridgeInfo	31
3.2.3. Consumer	31
3.2.4. ConsumerRecord	32
3.2.5. ConsumerRecordList	33
3.2.6. CreatedConsumer	33
3.2.7. Error	33
3.2.8. KafkaHeader	33
3.2.9. KafkaHeaderList	33
3.2.10. OffsetCommitSeek	33
3.2.11. OffsetCommitSeekList	34
3.2.12. OffsetRecordSent	34
3.2.13. OffsetRecordSentList	34
3.2.14. OffsetsSummary	34
3.2.15. Partition	35

3.2.16. PartitionMetadata	35
3.2.17. Partitions	35
3.2.18. ProducerRecord	35
3.2.19. ProducerRecordList	35
3.2.20. ProducerRecordToPartition	36
3.2.21. ProducerRecordToPartitionList	36
3.2.22. Replica	36
3.2.23. SubscribedTopicList	36
3.2.24. TopicMetadata	36
3.2.25. Topics	37
3.3. PATHS	37
3.3.1. GET /	37
3.3.1.1. Description	37
3.3.1.2. Responses	37
3.3.1.3. Produces	37
3.3.1.4. Example HTTP response	37
3.3.1.4.1. Response 200	37
3.3.2. POST /consumers/{groupid}	38
3.3.2.1. Description	38
3.3.2.2. Parameters	38
3.3.2.3. Responses	38
3.3.2.4. Consumes	38
3.3.2.5. Produces	38
3.3.2.6. Tags	38
3.3.2.7. Example HTTP request	39
3.3.2.7.1. Request body	39
3.3.2.8. Example HTTP response	39
3.3.2.8.1. Response 200	39
3.3.2.8.2. Response 409	39
3.3.2.8.3. Response 422	39
3.3.3. DELETE /consumers/{groupid}/instances/{name}	39
3.3.3.1. Description	39
3.3.3.2. Parameters	39
3.3.3.3. Responses	40
3.3.3.4. Consumes	40
3.3.3.5. Produces	40
3.3.3.6. Tags	40
3.3.3.7. Example HTTP response	40
3.3.3.7.1. Response 404	40
3.3.4. POST /consumers/{groupid}/instances/{name}/assignments	40
3.3.4.1. Description	40
3.3.4.2. Parameters	40
3.3.4.3. Responses	41
3.3.4.4. Consumes	41
3.3.4.5. Produces	41
3.3.4.6. Tags	41
3.3.4.7. Example HTTP request	41
3.3.4.7.1. Request body	41
3.3.4.8. Example HTTP response	42
3.3.4.8.1. Response 404	42
3.3.4.8.2. Response 409	42
3.3.5. POST /consumers/{groupid}/instances/{name}/offsets	42
3.3.5.1. Description	42

3.3.5.2. Parameters	42
3.3.5.3. Responses	42
3.3.5.4. Consumes	42
3.3.5.5. Produces	43
3.3.5.6. Tags	43
3.3.5.7. Example HTTP request	43
3.3.5.7.1. Request body	43
3.3.5.8. Example HTTP response	43
3.3.5.8.1. Response 404	43
3.3.6. POST /consumers/{groupid}/instances/{name}/positions	43
3.3.6.1. Description	43
3.3.6.2. Parameters	43
3.3.6.3. Responses	44
3.3.6.4. Consumes	44
3.3.6.5. Produces	44
3.3.6.6. Tags	44
3.3.6.7. Example HTTP request	44
3.3.6.7.1. Request body	44
3.3.6.8. Example HTTP response	45
3.3.6.8.1. Response 404	45
3.3.7. POST /consumers/{groupid}/instances/{name}/positions/beginning	45
3.3.7.1. Description	45
3.3.7.2. Parameters	45
3.3.7.3. Responses	45
3.3.7.4. Consumes	45
3.3.7.5. Produces	45
3.3.7.6. Tags	46
3.3.7.7. Example HTTP request	46
3.3.7.7.1. Request body	46
3.3.7.8. Example HTTP response	46
3.3.7.8.1. Response 404	46
3.3.8. POST /consumers/{groupid}/instances/{name}/positions/end	46
3.3.8.1. Description	46
3.3.8.2. Parameters	46
3.3.8.3. Responses	47
3.3.8.4. Consumes	47
3.3.8.5. Produces	47
3.3.8.6. Tags	47
3.3.8.7. Example HTTP request	47
3.3.8.7.1. Request body	47
3.3.8.8. Example HTTP response	47
3.3.8.8.1. Response 404	47
3.3.9. GET /consumers/{groupid}/instances/{name}/records	47
3.3.9.1. Description	48
3.3.9.2. Parameters	48
3.3.9.3. Responses	48
3.3.9.4. Produces	48
3.3.9.5. Tags	49
3.3.9.6. Example HTTP response	49
3.3.9.6.1. Response 200	49
3.3.9.6.2. Response 404	49
3.3.9.6.3. Response 406	49
3.3.9.6.4. Response 422	50

3.3.10. POST /consumers/{groupid}/instances/{name}/subscription	50
3.3.10.1. Description	50
3.3.10.2. Parameters	50
3.3.10.3. Responses	50
3.3.10.4. Consumes	50
3.3.10.5. Produces	51
3.3.10.6. Tags	51
3.3.10.7. Example HTTP request	51
3.3.10.7.1. Request body	51
3.3.10.8. Example HTTP response	51
3.3.10.8.1. Response 404	51
3.3.10.8.2. Response 409	51
3.3.10.8.3. Response 422	51
3.3.11. GET /consumers/{groupid}/instances/{name}/subscription	51
3.3.11.1. Description	51
3.3.11.2. Parameters	51
3.3.11.3. Responses	52
3.3.11.4. Produces	52
3.3.11.5. Tags	52
3.3.11.6. Example HTTP response	52
3.3.11.6.1. Response 200	52
3.3.11.6.2. Response 404	52
3.3.12. DELETE /consumers/{groupid}/instances/{name}/subscription	53
3.3.12.1. Description	53
3.3.12.2. Parameters	53
3.3.12.3. Responses	53
3.3.12.4. Tags	53
3.3.12.5. Example HTTP response	53
3.3.12.5.1. Response 404	53
3.3.13. GET /healthy	53
3.3.13.1. Description	53
3.3.13.2. Responses	53
3.3.14. GET /openapi	54
3.3.14.1. Description	54
3.3.14.2. Responses	54
3.3.14.3. Produces	54
3.3.15. GET /ready	54
3.3.15.1. Description	54
3.3.15.2. Responses	54
3.3.16. GET /topics	54
3.3.16.1. Description	54
3.3.16.2. Responses	54
3.3.16.3. Produces	55
3.3.16.4. Tags	55
3.3.16.5. Example HTTP response	55
3.3.16.5.1. Response 200	55
3.3.17. POST /topics/{topicname}	55
3.3.17.1. Description	55
3.3.17.2. Parameters	55
3.3.17.3. Responses	55
3.3.17.4. Consumes	56
3.3.17.5. Produces	56
3.3.17.6. Tags	56

3.3.17.7. Example HTTP request	56
3.3.17.7.1. Request body	56
3.3.17.8. Example HTTP response	56
3.3.17.8.1. Response 200	56
3.3.17.8.2. Response 404	56
3.3.17.8.3. Response 422	57
3.3.18. GET /topics/{topicname}	57
3.3.18.1. Description	57
3.3.18.2. Parameters	57
3.3.18.3. Responses	57
3.3.18.4. Produces	57
3.3.18.5. Tags	57
3.3.18.6. Example HTTP response	57
3.3.18.6.1. Response 200	57
3.3.19. GET /topics/{topicname}/partitions	58
3.3.19.1. Description	58
3.3.19.2. Parameters	58
3.3.19.3. Responses	58
3.3.19.4. Produces	59
3.3.19.5. Tags	59
3.3.19.6. Example HTTP response	59
3.3.19.6.1. Response 200	59
3.3.19.6.2. Response 404	59
3.3.20. POST /topics/{topicname}/partitions/{partitionid}	59
3.3.20.1. Description	59
3.3.20.2. Parameters	59
3.3.20.3. Responses	60
3.3.20.4. Consumes	60
3.3.20.5. Produces	60
3.3.20.6. Tags	60
3.3.20.7. Example HTTP request	60
3.3.20.7.1. Request body	60
3.3.20.8. Example HTTP response	61
3.3.20.8.1. Response 200	61
3.3.20.8.2. Response 404	61
3.3.20.8.3. Response 422	61
3.3.21. GET /topics/{topicname}/partitions/{partitionid}	61
3.3.21.1. Description	61
3.3.21.2. Parameters	61
3.3.21.3. Responses	62
3.3.21.4. Produces	62
3.3.21.5. Tags	62
3.3.21.6. Example HTTP response	62
3.3.21.6.1. Response 200	62
3.3.21.6.2. Response 404	62
3.3.22. GET /topics/{topicname}/partitions/{partitionid}/offsets	62
3.3.22.1. Description	63
3.3.22.2. Parameters	63
3.3.22.3. Responses	63
3.3.22.4. Produces	63
3.3.22.5. Tags	63
3.3.22.6. Example HTTP response	63
3.3.22.6.1. Response 200	63

3.3.22.6.2. Response 404	63
--------------------------	----

APPENDIX A. USING YOUR SUBSCRIPTION	64
Accessing Your Account	64
Activating a Subscription	64
Downloading Zip and Tar Files	64

MAKING OPEN SOURCE MORE INCLUSIVE

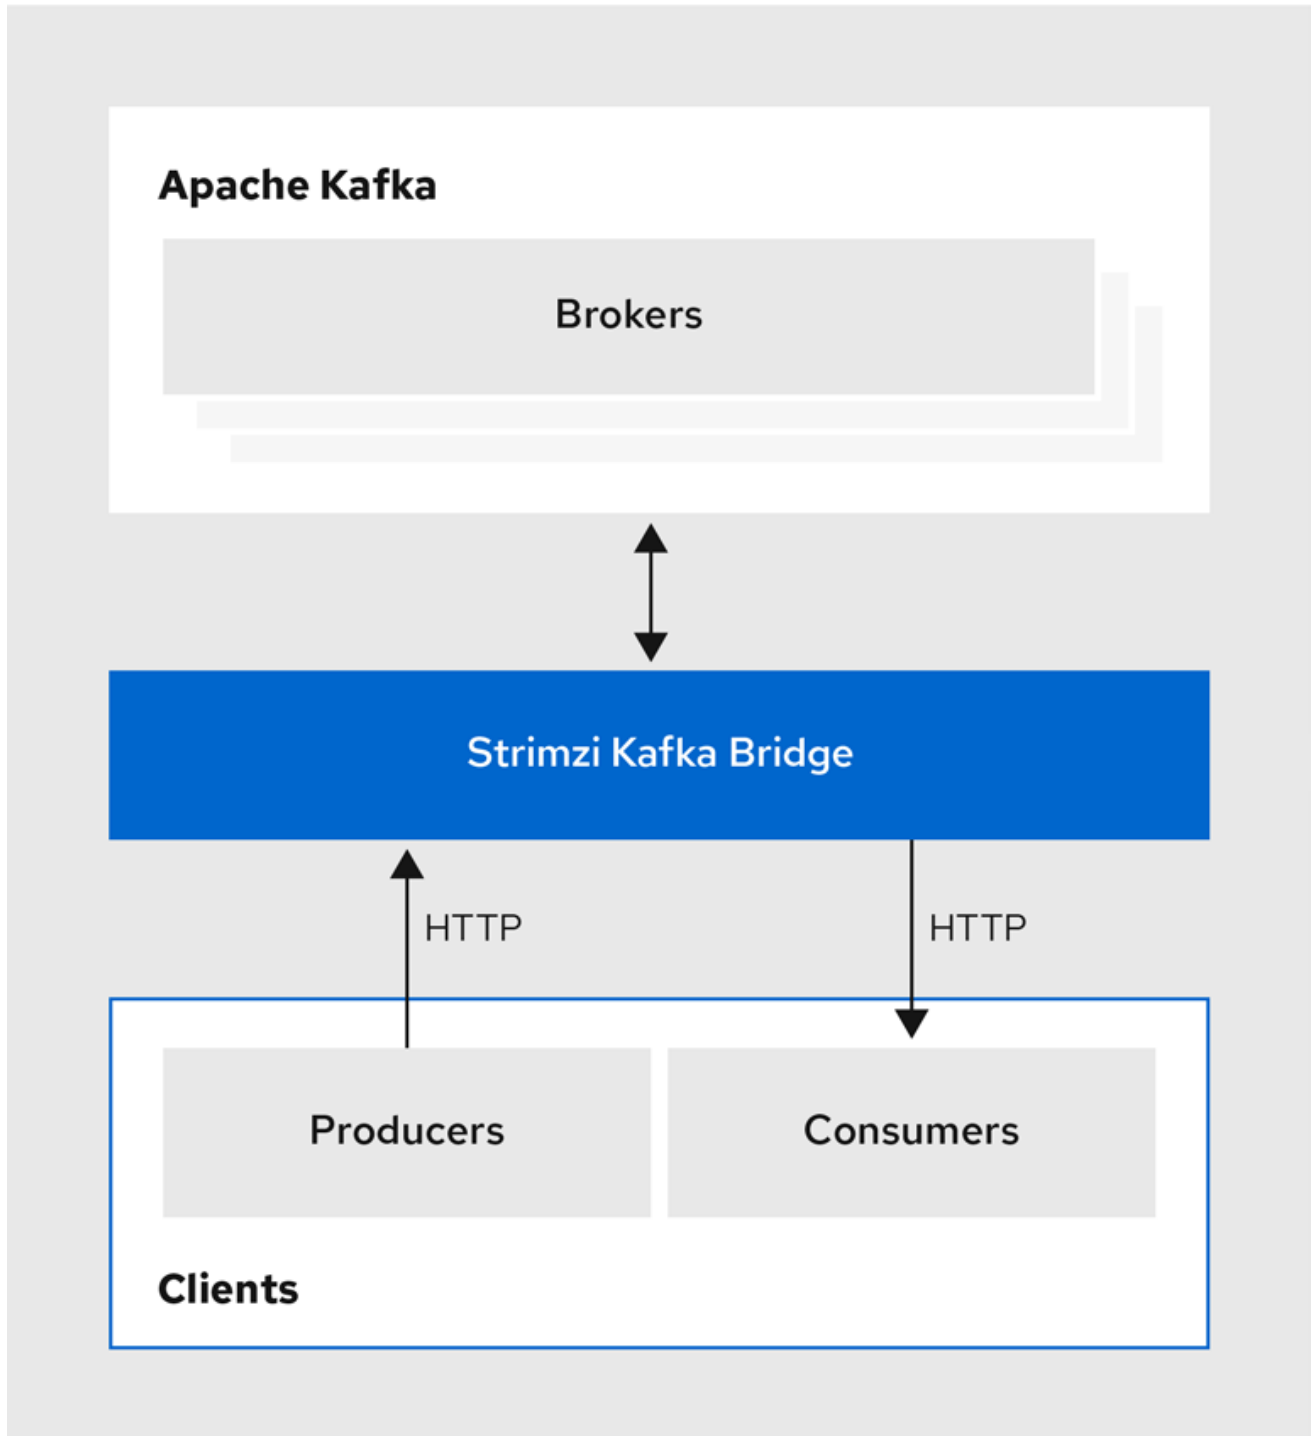
Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

CHAPTER 1. KAFKA BRIDGE OVERVIEW

Use the AMQ Streams Kafka Bridge to make HTTP requests to a Kafka cluster.

You can use the Kafka Bridge to integrate HTTP client applications with your Kafka cluster.

HTTP client integration



1.1. RUNNING THE KAFKA BRIDGE

Install the AMQ Streams Kafka Bridge to run in the same environment as your Kafka cluster.

You can download and add the Kafka Bridge installation artifacts to your host machine. To try out the Kafka Bridge in your local environment, see the [Kafka Bridge quickstart](#).

If you deployed AMQ Streams on OpenShift, you can use the AMQ Streams Cluster Operator to deploy the Kafka Bridge to the OpenShift cluster. You'll need a running Kafka cluster that was deployed by the Cluster Operator in an OpenShift namespace. You can configure your deployment to access the Kafka Bridge outside the OpenShift cluster.

Additional resources

- [AMQ Streams documentation](#)

1.2. KAFKA BRIDGE INTERFACE

The Kafka Bridge provides a RESTful interface that allows HTTP-based clients to interact with a Kafka cluster. It offers the advantages of a web API connection to AMQ Streams, without the need for client applications to interpret the Kafka protocol.

The API has two main resources – **consumers** and **topics** – that are exposed and made accessible through endpoints to interact with consumers and producers in your Kafka cluster. The resources relate only to the Kafka Bridge, not the consumers and producers connected directly to Kafka.

1.2.1. HTTP requests

The Kafka Bridge supports HTTP requests to a Kafka cluster, with methods to:

- Send messages to a topic.
- Retrieve messages from topics.
- Retrieve a list of partitions for a topic.
- Create and delete consumers.
- Subscribe consumers to topics, so that they start receiving messages from those topics.
- Retrieve a list of topics that a consumer is subscribed to.
- Unsubscribe consumers from topics.
- Assign partitions to consumers.
- Commit a list of consumer offsets.
- Seek on a partition, so that a consumer starts receiving messages from the first or last offset position, or a given offset position.

The methods provide JSON responses and HTTP response code error handling. Messages can be sent in JSON or binary formats.

Clients can produce and consume messages without the requirement to use the native Kafka protocol.

Additional resources

- [AMQ Streams Kafka Bridge API reference](#)

1.3. SECURING CONNECTIVITY TO THE KAFKA CLUSTER

You can configure the following between the Kafka Bridge and your Kafka cluster:

- TLS or SASL-based authentication
- A TLS-encrypted connection

You configure the Kafka Bridge for authentication through its [properties file](#).

You can also use ACLs in Kafka brokers to restrict the topics that can be consumed and produced using the Kafka Bridge.

Additional resources

- [AMQ Streams documentation](#)

1.4. SECURING THE KAFKA BRIDGE HTTP INTERFACE

Authentication and encryption between HTTP clients and the Kafka Bridge is not supported directly by the Kafka Bridge. Requests sent from clients to the Kafka Bridge are sent without authentication or encryption. Requests must use HTTP rather than HTTPS.

You can combine the Kafka Bridge with the following tools to secure it:

- Network policies and firewalls that define which pods can access the Kafka Bridge
- Reverse proxies (for example, OAuth 2.0)
- API gateways

1.5. REQUESTS TO THE KAFKA BRIDGE

Specify data formats and HTTP headers to ensure valid requests are submitted to the Kafka Bridge.

1.5.1. Content Type headers

API request and response bodies are always encoded as JSON.

- When performing consumer operations, **POST** requests must provide the following **Content-Type** header if there is a non-empty body:

Content-Type: application/vnd.kafka.v2+json

- When performing producer operations, **POST** requests must provide **Content-Type** headers specifying the *embedded data format* of the messages produced. This can be either **json** or **binary**.

Embedded data format	Content-Type header
JSON	Content-Type: application/vnd.kafka.json.v2+json

Embedded data format	Content-Type header
Binary	Content-Type: application/vnd.kafka.binary.v2+json

The embedded data format is set per consumer, as described in the next section.

The **Content-Type** must *not* be set if the **POST** request has an empty body. An empty body can be used to create a consumer with the default values.

1.5.2. Embedded data format

The embedded data format is the format of the Kafka messages that are transmitted, over HTTP, from a producer to a consumer using the Kafka Bridge. Two embedded data formats are supported: JSON and binary.

When creating a consumer using the `/consumers/groupid` endpoint, the **POST** request body must specify an embedded data format of either JSON or binary. This is specified in the **format** field, for example:

```
{
  "name": "my-consumer",
  "format": "binary", 1
  # ...
}
```

1 A binary embedded data format.

The embedded data format specified when creating a consumer must match the data format of the Kafka messages it will consume.

If you choose to specify a binary embedded data format, subsequent producer requests must provide the binary data in the request body as Base64-encoded strings. For example, when sending messages using the `/topics/topicname` endpoint, **records.value** must be encoded in Base64:

```
{
  "records": [
    {
      "key": "my-key",
      "value": "ZWR3YXJkdGhldGhyZWVsZWdnZWVjYXQ="
    },
  ],
}
```

Producer requests must also provide a **Content-Type** header that corresponds to the embedded data format, for example, **Content-Type: application/vnd.kafka.binary.v2+json**.

1.5.3. Message format

When sending messages using the `/topics` endpoint, you enter the message payload in the request body, in the **records** parameter.

The **records** parameter can contain any of these optional fields:

- Message **headers**
- Message **key**
- Message **value**
- Destination **partition**

Example POST request to /topics

```
curl -X POST \
  http://localhost:8080/topics/my-topic \
  -H 'content-type: application/vnd.kafka.json.v2+json' \
  -d '{
    "records": [
      {
        "key": "my-key",
        "value": "sales-lead-0001"
        "partition": 2
        "headers": [
          {
            "key": "key1",
            "value": "QXBhY2hIIEthZmthIGlzIHRoZSBib21iIQ==" 1
          }
        ]
      },
    ]
  }'
```

- 1 The header value in binary format and encoded as Base64.

1.5.4. Accept headers

After creating a consumer, all subsequent GET requests must provide an **Accept** header in the following format:

```
Accept: application/vnd.kafka.EMBEDDED-DATA-FORMAT.v2+json
```

The **EMBEDDED-DATA-FORMAT** is either **json** or **binary**.

For example, when retrieving records for a subscribed consumer using an embedded data format of JSON, include this Accept header:

```
Accept: application/vnd.kafka.json.v2+json
```

1.6. CORS

Cross-Origin Resource Sharing (CORS) allows you to specify allowed methods and originating URLs for accessing the Kafka cluster in your Kafka Bridge HTTP configuration.

Example CORS configuration for Kafka Bridge

```
# ...
http.cors.enabled=true
http.cors.allowedOrigins=https://strimzi.io
http.cors.allowedMethods=GET,POST,PUT,DELETE,OPTIONS,PATCH
```

CORS allows for *simple* and *preflighted* requests between origin sources on different domains.

Simple requests are suitable for standard requests using **GET**, **HEAD**, **POST** methods.

A preflighted request sends a *HTTP OPTIONS* request as an initial check that the actual request is safe to send. On confirmation, the actual request is sent. Preflight requests are suitable for methods that require greater safeguards, such as **PUT** and **DELETE**, and use non-standard headers.

All requests require an *origins* value in their header, which is the source of the HTTP request.

1.6.1. Simple request

For example, this simple request header specifies the origin as **https://strimzi.io**.

```
Origin: https://strimzi.io
```

The header information is added to the request.

```
curl -v -X GET HTTP-ADDRESS/bridge-consumer/records \
-H 'Origin: https://strimzi.io\'
-H 'content-type: application/vnd.kafka.v2+json'
```

In the response from the Kafka Bridge, an **Access-Control-Allow-Origin** header is returned.

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: * 1
```

1 Returning an asterisk (*) shows the resource can be accessed by any domain.

1.6.2. Preflighted request

An initial preflight request is sent to Kafka Bridge using an **OPTIONS** method. The *HTTP OPTIONS* request sends header information to check that Kafka Bridge will allow the actual request.

Here the preflight request checks that a **POST** request is valid from **https://strimzi.io**.

```
OPTIONS /my-group/instances/my-user/subscription HTTP/1.1
Origin: https://strimzi.io
Access-Control-Request-Method: POST 1
Access-Control-Request-Headers: Content-Type 2
```

1 Kafka Bridge is alerted that the actual request is a **POST** request.

2 The actual request will be sent with a **Content-Type** header.

OPTIONS is added to the header information of the preflight request.

```
curl -v -X OPTIONS -H 'Origin: https://strimzi.io' \
-H 'Access-Control-Request-Method: POST' \
-H 'content-type: application/vnd.kafka.v2+json'
```

Kafka Bridge responds to the initial request to confirm that the request will be accepted. The response header returns allowed origins, methods and headers.

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: https://strimzi.io
Access-Control-Allow-Methods: GET,POST,PUT,DELETE,OPTIONS,PATCH
Access-Control-Allow-Headers: content-type
```

If the origin or method is rejected, an error message is returned.

The actual request does not require **Access-Control-Request-Method** header, as it was confirmed in the preflight request, but it does require the origin header.

```
curl -v -X POST HTTP-ADDRESS/topics/bridge-topic \
-H 'Origin: https://strimzi.io' \
-H 'content-type: application/vnd.kafka.v2+json'
```

The response shows the originating URL is allowed.

```
HTTP/1.1 200 OK
Access-Control-Allow-Origin: https://strimzi.io
```

Additional resources

- [Fetch CORS specification](#)

1.7. CONFIGURING LOGGERS FOR THE KAFKA BRIDGE

The AMQ Streams Kafka bridge allows you to set a different log level for each operation that is defined by the related OpenAPI specification.

Each operation has a corresponding API endpoint through which the bridge receives requests from HTTP clients. You can change the log level on each endpoint to produce more or less fine-grained logging information about the incoming and outgoing HTTP requests.

Loggers are defined in the **log4j.properties** file, which has the following default configuration for **healthy** and **ready** endpoints:

```
log4j.logger.http.openapi.operation.healthy=WARN, out
log4j.additivity.http.openapi.operation.healthy=false
log4j.logger.http.openapi.operation.ready=WARN, out
log4j.additivity.http.openapi.operation.ready=false
```

The log level of all other operations is set to **INFO** by default. Loggers are formatted as follows:

```
log4j.logger.http.openapi.operation.<operation_id>
```

Where **<operation_id>** is the identifier of the specific operation.

List of operations defined by the OpenAPI specification

- **createConsumer**
- **deleteConsumer**
- **subscribe**
- **unsubscribe**
- **poll**
- **assign**
- **commit**
- **send**
- **sendToPartition**
- **seekToBeginning**
- **seekToEnd**
- **seek**
- **healthy**
- **ready**
- **openapi**

CHAPTER 2. KAFKA BRIDGE QUICKSTART

Use this quickstart to try out the AMQ Streams Kafka Bridge in your local development environment.

You will learn how to do the following:

- Produce messages to topics and partitions in your Kafka cluster
- Create a Kafka Bridge consumer
- Perform basic consumer operations, such as subscribing the consumer to topics and retrieving the messages that you produced

In this quickstart, HTTP requests are formatted as curl commands that you can copy and paste to your terminal.

Ensure you have the prerequisites and then follow the tasks in the order provided in this chapter.

About data formats

In this quickstart, you will produce and consume messages in JSON format, not binary.

Prerequisites for the quickstart

- A Kafka cluster is running on the host machine.

2.1. DOWNLOADING A KAFKA BRIDGE ARCHIVE

A zipped distribution of the AMQ Streams Kafka Bridge is available for download.

Procedure

- Download the latest version of the AMQ Streams Kafka Bridge archive from the [Customer Portal](#).

2.2. CONFIGURING KAFKA BRIDGE PROPERTIES

This procedure describes how to configure the Kafka and HTTP connection properties used by the AMQ Streams Kafka Bridge.

You configure the Kafka Bridge, as any other Kafka client, using appropriate prefixes for Kafka-related properties.

- **kafka.** for general configuration that applies to producers and consumers, such as server connection and security.
- **kafka.consumer.** for consumer-specific configuration passed only to the consumer.
- **kafka.producer.** for producer-specific configuration passed only to the producer.

As well as enabling HTTP access to a Kafka cluster, HTTP properties provide the capability to enable and define access control for the Kafka Bridge through Cross-Origin Resource Sharing (CORS). CORS is a HTTP mechanism that allows browser access to selected resources from more than one origin. To

configure CORS, you define a list of allowed resource origins and HTTP methods to access them. Additional HTTP headers in requests describe the CORS origins that are permitted access to the Kafka cluster.

Prerequisites

- [The Kafka Bridge installation archive is downloaded](#)

Procedure

1. Edit the **application.properties** file provided with the AMQ Streams Kafka Bridge installation archive.
Use the properties file to specify Kafka and HTTP-related properties, and to enable distributed tracing.
 - a. Configure standard Kafka-related properties, including properties specific to the Kafka consumers and producers.
Use:
 - **kafka.bootstrap.servers** to define the host/port connections to the Kafka cluster
 - **kafka.producer.acks** to provide acknowledgments to the HTTP client
 - **kafka.consumer.auto.offset.reset** to determine how to manage reset of the offset in Kafka
For more information on configuration of Kafka properties, see the [Apache Kafka website](#)
 - b. Configure HTTP-related properties to enable HTTP access to the Kafka cluster.
For example:

```
bridge.id=my-bridge
http.enabled=true
http.host=0.0.0.0
http.port=8080 1
http.cors.enabled=true 2
http.cors.allowedOrigins=https://strimzi.io 3
http.cors.allowedMethods=GET,POST,PUT,DELETE,OPTIONS,PATCH 4
```

- 1 The default HTTP configuration for the Kafka Bridge to listen on port 8080.
- 2 Set to **true** to enable CORS.
- 3 Comma-separated list of allowed CORS origins. You can use a URL or a Java regular expression.
- 4 Comma-separated list of allowed HTTP methods for CORS.

2.3. INSTALLING THE KAFKA BRIDGE

Follow this procedure to install the AMQ Streams Kafka Bridge.

Prerequisites

- [The Kafka Bridge installation archive is downloaded](#)
- [The Kafka Bridge configuration properties are set](#)

Procedure

1. If you have not already done so, unzip the Kafka Bridge installation archive to any directory.
2. Run the Kafka Bridge script using the configuration properties as a parameter:
For example:

```
./bin/kafka_bridge_run.sh --config-file=<path>/configfile.properties
```

3. Check to see that the installation was successful in the log.

```
HTTP-Kafka Bridge started and listening on port 8080
HTTP-Kafka Bridge bootstrap servers localhost:9092
```

2.4. PRODUCING MESSAGES TO TOPICS AND PARTITIONS

Use the Kafka Bridge to produce messages to a Kafka topic in JSON format by using the topics endpoint.

You can produce messages to topics in JSON format by using the [topics](#) endpoint. You can specify destination partitions for messages in the request body. The [partitions](#) endpoint provides an alternative method for specifying a single destination partition for all messages as a path parameter.

In this procedure, messages are produced to a topic called **bridge-quickstart-topic**.

Prerequisites


- The Kafka cluster has a topic with three partitions.
You can use the **kafka-topics.sh** utility to create topics.

Example topic creation with three partitions

```
bin/kafka-topics.sh --bootstrap-server localhost:9092 --create --topic bridge-quickstart-topic -
-partitions 3 --replication-factor 1
```

Verifying the topic was created

```
bin/kafka-topics.sh --bootstrap-server localhost:9092 --describe --topic bridge-quickstart-
topic
```



NOTE

If you deployed AMQ Streams on OpenShift, you can create a topic using the **KafkaTopic** custom resource.

Procedure

1. Using the Kafka Bridge, produce three messages to the topic you created:

-

```

curl -X POST \
  http://localhost:8080/topics/bridge-quickstart-topic \
  -H 'content-type: application/vnd.kafka.json.v2+json' \
  -d '{
    "records": [
      {
        "key": "my-key",
        "value": "sales-lead-0001"
      },
      {
        "value": "sales-lead-0002",
        "partition": 2
      },
      {
        "value": "sales-lead-0003"
      }
    ]
  }'
```

- **sales-lead-0001** is sent to a partition based on the hash of the key.
 - **sales-lead-0002** is sent directly to partition 2.
 - **sales-lead-0003** is sent to a partition in the **bridge-quickstart-topic** topic using a round-robin method.
2. If the request is successful, the Kafka Bridge returns an **offsets** array, along with a **200** code and a **content-type** header of **application/vnd.kafka.v2+json**. For each message, the **offsets** array describes:
- The partition that the message was sent to
 - The current message offset of the partition

Example response

```

#...
{
  "offsets":[
    {
      "partition":0,
      "offset":0
    },
    {
      "partition":2,
      "offset":0
    },
    {
      "partition":0,
      "offset":1
    }
  ]
}
```

Additional topic requests

Make other curl requests to find information on topics and partitions.

List topics

```
curl -X GET \
  http://localhost:8080/topics
```

Example response

```
[
  "__strimzi_store_topic",
  "__strimzi-topic-operator-kstreams-topic-store-changelog",
  "bridge-quickstart-topic",
  "my-topic"
]
```

Get topic configuration and partition details

```
curl -X GET \
  http://localhost:8080/topics/bridge-quickstart-topic
```

Example response

```
{
  "name": "bridge-quickstart-topic",
  "configs": {
    "compression.type": "producer",
    "leader.replication.throttled.replicas": "",
    "min.insync.replicas": "1",
    "message.downconversion.enable": "true",
    "segment.jitter.ms": "0",
    "cleanup.policy": "delete",
    "flush.ms": "9223372036854775807",
    "follower.replication.throttled.replicas": "",
    "segment.bytes": "1073741824",
    "retention.ms": "604800000",
    "flush.messages": "9223372036854775807",
    "message.format.version": "2.8-IV1",
    "max.compaction.lag.ms": "9223372036854775807",
    "file.delete.delay.ms": "60000",
    "max.message.bytes": "1048588",
    "min.compaction.lag.ms": "0",
    "message.timestamp.type": "CreateTime",
    "preallocate": "false",
    "index.interval.bytes": "4096",
    "min.cleanable.dirty.ratio": "0.5",
    "unclean.leader.election.enable": "false",
    "retention.bytes": "-1",
    "delete.retention.ms": "86400000",
    "segment.ms": "604800000",
    "message.timestamp.difference.max.ms": "9223372036854775807",
    "segment.index.bytes": "10485760"
  },
  "partitions": [
```

```
{
  "partition": 0,
  "leader": 0,
  "replicas": [
    {
      "broker": 0,
      "leader": true,
      "in_sync": true
    },
    {
      "broker": 1,
      "leader": false,
      "in_sync": true
    },
    {
      "broker": 2,
      "leader": false,
      "in_sync": true
    }
  ]
},
{
  "partition": 1,
  "leader": 2,
  "replicas": [
    {
      "broker": 2,
      "leader": true,
      "in_sync": true
    },
    {
      "broker": 0,
      "leader": false,
      "in_sync": true
    },
    {
      "broker": 1,
      "leader": false,
      "in_sync": true
    }
  ]
},
{
  "partition": 2,
  "leader": 1,
  "replicas": [
    {
      "broker": 1,
      "leader": true,
      "in_sync": true
    },
    {
      "broker": 2,
      "leader": false,
      "in_sync": true
    }
  ]
},
```

```

    {
      "broker": 0,
      "leader": false,
      "in_sync": true
    }
  ]
}
]
}

```

List the partitions of a specific topic

```

curl -X GET \
  http://localhost:8080/topics/bridge-quickstart-topic/partitions

```

Example response

```

[
  {
    "partition": 0,
    "leader": 0,
    "replicas": [
      {
        "broker": 0,
        "leader": true,
        "in_sync": true
      },
      {
        "broker": 1,
        "leader": false,
        "in_sync": true
      },
      {
        "broker": 2,
        "leader": false,
        "in_sync": true
      }
    ]
  },
  {
    "partition": 1,
    "leader": 2,
    "replicas": [
      {
        "broker": 2,
        "leader": true,
        "in_sync": true
      },
      {
        "broker": 0,
        "leader": false,
        "in_sync": true
      },
      {
        "broker": 1,

```

```
    "leader": false,
    "in_sync": true
  }
]
},
{
  "partition": 2,
  "leader": 1,
  "replicas": [
    {
      "broker": 1,
      "leader": true,
      "in_sync": true
    },
    {
      "broker": 2,
      "leader": false,
      "in_sync": true
    },
    {
      "broker": 0,
      "leader": false,
      "in_sync": true
    }
  ]
}
]
```

List the details of a specific topic partition

```
curl -X GET \
  http://localhost:8080/topics/bridge-quickstart-topic/partitions/0
```

Example response

```
{
  "partition": 0,
  "leader": 0,
  "replicas": [
    {
      "broker": 0,
      "leader": true,
      "in_sync": true
    },
    {
      "broker": 1,
      "leader": false,
      "in_sync": true
    },
    {
      "broker": 2,
      "leader": false,
      "in_sync": true
    }
  ]
}
```

```
}
]
```

List the offsets of a specific topic partition

```
curl -X GET \
  http://localhost:8080/topics/bridge-quickstart-topic/partitions/0/offsets
```

Example response

```
{
  "beginning_offset": 0,
  "end_offset": 1
}
```

What to do next

After producing messages to topics and partitions, [create a Kafka Bridge consumer](#) .

Additional resources

- [POST /topics/{topicname}](#)
- [POST /topics/{topicname}/partitions/{partitionid}](#)

2.5. CREATING A KAFKA BRIDGE CONSUMER

Before you can perform any consumer operations in the Kafka cluster, you must first create a consumer by using the [consumers](#) endpoint. The consumer is referred to as a *Kafka Bridge consumer*.

Procedure

1. Create a Kafka Bridge consumer in a new consumer group named **bridge-quickstart-consumer-group**:

```
curl -X POST http://localhost:8080/consumers/bridge-quickstart-consumer-group \
  -H 'content-type: application/vnd.kafka.v2+json' \
  -d '{
    "name": "bridge-quickstart-consumer",
    "auto.offset.reset": "earliest",
    "format": "json",
    "enable.auto.commit": false,
    "fetch.min.bytes": 512,
    "consumer.request.timeout.ms": 30000
  }'
```

- The consumer is named **bridge-quickstart-consumer** and the embedded data format is set as **json**.
- Some basic configuration settings are defined.

- The consumer will not commit offsets to the log automatically because the **enable.auto.commit** setting is **false**. You will commit the offsets manually later in this quickstart.
If the request is successful, the Kafka Bridge returns the consumer ID (**instance_id**) and base URL (**base_uri**) in the response body, along with a **200** code.

Example response

```
#...
{
  "instance_id": "bridge-quickstart-consumer",
  "base_uri": "http://<bridge_id>-bridge-service:8080/consumers/bridge-quickstart-consumer-group/instances/bridge-quickstart-consumer"
}
```

2. Copy the base URL (**base_uri**) to use in the other consumer operations in this quickstart.

What to do next

Now that you have created a Kafka Bridge consumer, you can [subscribe it to topics](#).

Additional resources

- [POST /consumers/{groupid}](#)

2.6. SUBSCRIBING A KAFKA BRIDGE CONSUMER TO TOPICS

After you have created a Kafka Bridge consumer, subscribe it to one or more topics by using the [subscription](#) endpoint. When subscribed, the consumer starts receiving all messages that are produced to the topic.

Procedure

- Subscribe the consumer to the **bridge-quickstart-topic** topic that you created earlier, in [Producing messages to topics and partitions](#) :

```
curl -X POST http://localhost:8080/consumers/bridge-quickstart-consumer-group/instances/bridge-quickstart-consumer/subscription \
-H 'content-type: application/vnd.kafka.v2+json' \
-d '{
  "topics": [
    "bridge-quickstart-topic"
  ]
}'
```

The **topics** array can contain a single topic (as shown here) or multiple topics. If you want to subscribe the consumer to multiple topics that match a regular expression, you can use the **topic_pattern** string instead of the **topics** array.

If the request is successful, the Kafka Bridge returns a **204** (No Content) code only.

What to do next

After subscribing a Kafka Bridge consumer to topics, you can [retrieve messages from the consumer](#).

Additional resources

Additional resources

- [POST /consumers/{groupid}/instances/{name}/subscription](#)

2.7. RETRIEVING THE LATEST MESSAGES FROM A KAFKA BRIDGE CONSUMER

Retrieve the latest messages from the Kafka Bridge consumer by requesting data from the [records](#) endpoint. In production, HTTP clients can call this endpoint repeatedly (in a loop).

Procedure

1. Produce additional messages to the Kafka Bridge consumer, as described in [Producing messages to topics and partitions](#).
2. Submit a **GET** request to the **records** endpoint:

```
curl -X GET http://localhost:8080/consumers/bridge-quickstart-consumer-
group/instances/bridge-quickstart-consumer/records \
-H 'accept: application/vnd.kafka.json.v2+json'
```

After creating and subscribing to a Kafka Bridge consumer, a first GET request will return an empty response because the poll operation starts a rebalancing process to assign partitions.

3. Repeat step two to retrieve messages from the Kafka Bridge consumer. The Kafka Bridge returns an array of messages – describing the topic name, key, value, partition, and offset – in the response body, along with a **200** code. Messages are retrieved from the latest offset by default.

```
HTTP/1.1 200 OK
content-type: application/vnd.kafka.json.v2+json
#...
[
  {
    "topic":"bridge-quickstart-topic",
    "key":"my-key",
    "value":"sales-lead-0001",
    "partition":0,
    "offset":0
  },
  {
    "topic":"bridge-quickstart-topic",
    "key":null,
    "value":"sales-lead-0003",
    "partition":0,
    "offset":1
  },
  #...
```

**NOTE**

If an empty response is returned, produce more records to the consumer as described in [Producing messages to topics and partitions](#), and then try retrieving messages again.

What to do next

After retrieving messages from a Kafka Bridge consumer, try [committing offsets to the log](#).

Additional resources

- [GET /consumers/{groupid}/instances/{name}/records](#)

2.8. COMMITTING OFFSETS TO THE LOG

Use the [offsets](#) endpoint to manually commit offsets to the log for all messages received by the Kafka Bridge consumer. This is required because the Kafka Bridge consumer that you created earlier, in [Creating a Kafka Bridge consumer](#), was configured with the `enable.auto.commit` setting as `false`.

Procedure

- Commit offsets to the log for the **bridge-quickstart-consumer**:

```
curl -X POST http://localhost:8080/consumers/bridge-quickstart-consumer-
group/instances/bridge-quickstart-consumer/offsets
```

Because no request body is submitted, offsets are committed for all the records that have been received by the consumer. Alternatively, the request body can contain an array ([OffsetCommitSeekList](#)) that specifies the topics and partitions that you want to commit offsets for.

If the request is successful, the Kafka Bridge returns a **204** code only.

What to do next

After committing offsets to the log, try out the endpoints for [seeking to offsets](#).

Additional resources

- [POST /consumers/{groupid}/instances/{name}/offsets](#)

2.9. SEEKING TO OFFSETS FOR A PARTITION

Use the [positions](#) endpoints to configure the Kafka Bridge consumer to retrieve messages for a partition from a specific offset, and then from the latest offset. This is referred to in Apache Kafka as a seek operation.

Procedure

1. Seek to a specific offset for partition 0 of the **quickstart-bridge-topic** topic:

```
curl -X POST http://localhost:8080/consumers/bridge-quickstart-consumer-
group/instances/bridge-quickstart-consumer/positions \
-H 'content-type: application/vnd.kafka.v2+json' \
-d '{
  "offsets": [
    {
      "topic": "bridge-quickstart-topic",
      "partition": 0,
      "offset": 2
    }
  ]
}
```



```
}
  ]
}'
```

If the request is successful, the Kafka Bridge returns a **204** code only.

2. Submit a **GET** request to the **records** endpoint:


```
curl -X GET http://localhost:8080/consumers/bridge-quickstart-consumer-
group/instances/bridge-quickstart-consumer/records \
-H 'accept: application/vnd.kafka.json.v2+json'
```

The Kafka Bridge returns messages from the offset that you seeked to.

3. Restore the default message retrieval behavior by seeking to the last offset for the same partition. This time, use the [positions/end](#) endpoint.

```
curl -X POST http://localhost:8080/consumers/bridge-quickstart-consumer-
group/instances/bridge-quickstart-consumer/positions/end \
-H 'content-type: application/vnd.kafka.v2+json' \
-d '{
  "partitions": [
    {
      "topic": "bridge-quickstart-topic",
      "partition": 0
    }
  ]
}'
```

If the request is successful, the Kafka Bridge returns another **204** code.



NOTE

You can also use the [positions/beginning](#) endpoint to seek to the first offset for one or more partitions.

What to do next

In this quickstart, you have used the AMQ Streams Kafka Bridge to perform several common operations on a Kafka cluster. You can now [delete the Kafka Bridge consumer](#) that you created earlier.

Additional resources

- [POST /consumers/{groupid}/instances/{name}/positions](#)
- [POST /consumers/{groupid}/instances/{name}/positions/beginning](#)
- [POST /consumers/{groupid}/instances/{name}/positions/end](#)

2.10. DELETING A KAFKA BRIDGE CONSUMER

Delete the Kafka Bridge consumer that you used throughout this quickstart.

Procedure

- Delete the Kafka Bridge consumer by sending a **DELETE** request to the [instances](#) endpoint.

```
curl -X DELETE http://localhost:8080/consumers/bridge-quickstart-consumer-  
group/instances/bridge-quickstart-consumer
```

If the request is successful, the Kafka Bridge returns a **204** code.

Additional resources

- [DELETE /consumers/{groupid}/instances/{name}](#)

CHAPTER 3. AMQ STREAMS KAFKA BRIDGE API REFERENCE

3.1. OVERVIEW

The AMQ Streams Kafka Bridge provides a REST API for integrating HTTP based client applications with a Kafka cluster. You can use the API to create and manage consumers and send and receive records over HTTP rather than the native Kafka protocol.

3.1.1. Version information

Version : 0.1.0

3.1.2. Tags

- **Consumers** : Consumer operations to create consumers in your Kafka cluster and perform common actions, such as subscribing to topics, retrieving processed records, and committing offsets.
- **Producer** : Producer operations to send records to a specified topic or topic partition.
- **Seek** : Seek operations that enable a consumer to begin receiving messages from a given offset position.
- **Topics** : Topic operations to send messages to a specified topic or topic partition, optionally including message keys in requests. You can also retrieve topics and topic metadata.

3.1.3. Consumes

- **application/json**

3.1.4. Produces

- **application/json**

3.2. DEFINITIONS

3.2.1. AssignedTopicPartitions

Type : < string, < integer (int32) > array > map

3.2.2. BridgeInfo

Information about Kafka Bridge instance.

Name	Schema
<code>bridge_version</code> <i>optional</i>	string

3.2.3. Consumer

Name	Description	Schema
auto.offset.reset <i>optional</i>	Resets the offset position for the consumer. If set to earliest , messages are read from the first offset. If set to latest , messages are read from the latest offset.	string
consumer.request.timeout.ms <i>optional</i>	Sets the maximum amount of time, in milliseconds, for the consumer to wait for messages for a request. If the timeout period is reached without a response, an error is returned.	integer
enable.auto.commit <i>optional</i>	If set to true , message offsets are committed automatically for the consumer. If set to false , message offsets must be committed manually.	boolean
fetch.min.bytes <i>optional</i>	Sets the minimum amount of data, in bytes, for the consumer to receive. The broker waits until the data to send exceeds this amount.	integer
format <i>optional</i>	The allowable message format for the consumer, which can be binary (default) or json . The messages are converted into a JSON format.	string
isolation.level <i>optional</i>	If set to read_uncommitted , all transaction records are retrieved. If set to read_committed , the records from committed transactions are retrieved.	string
name <i>optional</i>	The unique name for the consumer instance. The name is unique within the scope of the consumer group. The name is used in URLs.	string

3.2.4. ConsumerRecord

Name	Schema
headers <i>optional</i>	KafkaHeaderList
key <i>optional</i>	string
offset <i>optional</i>	integer (int64)
partition <i>optional</i>	integer (int32)
topic <i>optional</i>	string

Name	Schema
value <i>optional</i>	string

3.2.5. ConsumerRecordList

Type : < [ConsumerRecord](#) > array

3.2.6. CreatedConsumer

Name	Description	Schema
base_uri <i>optional</i>	Base URI used to construct URIs for subsequent requests against this consumer instance.	string
instance_id <i>optional</i>	Unique ID for the consumer instance in the group.	string

3.2.7. Error

Name	Schema
error_code <i>optional</i>	integer (int32)
message <i>optional</i>	string

3.2.8. KafkaHeader

Name	Description	Schema
key <i>required</i>		string
value <i>required</i>	The header value in binary format, base64-encoded Pattern : <code>""^(?:[A-Za-z0-9+]{4})*(?:[A-Za-z0-9+]{2}=[A-Za-z0-9+]{3}=)?\$"</code>	string (byte)

3.2.9. KafkaHeaderList

Type : < [KafkaHeader](#) > array

3.2.10. OffsetCommitSeek

Name	Schema
offset <i>required</i>	integer (int64)
partition <i>required</i>	integer (int32)
topic <i>required</i>	string

3.2.11. OffsetCommitSeekList

Name	Schema
offsets <i>optional</i>	< OffsetCommitSeek > array

3.2.12. OffsetRecordSent

Name	Schema
offset <i>optional</i>	integer (int64)
partition <i>optional</i>	integer (int32)

3.2.13. OffsetRecordSentList

Name	Schema
offsets <i>optional</i>	< OffsetRecordSent > array

3.2.14. OffsetsSummary

Name	Schema
beginning_offset <i>optional</i>	integer (int64)
end_offset <i>optional</i>	integer (int64)

3.2.15. Partition

Name	Schema
partition <i>optional</i>	integer (int32)
topic <i>optional</i>	string

3.2.16. PartitionMetadata

Name	Schema
leader <i>optional</i>	integer (int32)
partition <i>optional</i>	integer (int32)
replicas <i>optional</i>	< Replica > array

3.2.17. Partitions

Name	Schema
partitions <i>optional</i>	< Partition > array

3.2.18. ProducerRecord

Name	Schema
headers <i>optional</i>	KafkaHeaderList
partition <i>optional</i>	integer (int32)

3.2.19. ProducerRecordList

Name	Schema
records <i>optional</i>	< ProducerRecord > array

3.2.20. ProducerRecordToPartition

Type: object

3.2.21. ProducerRecordToPartitionList

Name	Schema
records <i>optional</i>	< ProducerRecordToPartition > array

3.2.22. Replica

Name	Schema
broker <i>optional</i>	integer (int32)
in_sync <i>optional</i>	boolean
leader <i>optional</i>	boolean

3.2.23. SubscribedTopicList

Name	Schema
partitions <i>optional</i>	< AssignedTopicPartitions > array
topics <i>optional</i>	Topics

3.2.24. TopicMetadata

Name	Description	Schema
configs <i>optional</i>	Per-topic configuration overrides	< string, string > map
name <i>optional</i>	Name of the topic	string
partitions <i>optional</i>		< PartitionMetadata > array

3.2.25. Topics

Name	Description	Schema
topic_pattern <i>optional</i>	A regex topic pattern for matching multiple topics	string
topics <i>optional</i>		< string > array

3.3. PATHS

3.3.1. GET /

3.3.1.1. Description

Retrieves information about the Kafka Bridge instance, in JSON format.

3.3.1.2. Responses

HTTP Code	Description	Schema
200	Information about Kafka Bridge instance.	BridgeInfo

3.3.1.3. Produces

- **application/json**

3.3.1.4. Example HTTP response

3.3.1.4.1. Response 200

```
{
  "bridge_version" : "0.16.0"
}
```

3.3.2. POST /consumers/{groupid}

3.3.2.1. Description

Creates a consumer instance in the given consumer group. You can optionally specify a consumer name and supported configuration options. It returns a base URI which must be used to construct URLs for subsequent requests against this consumer instance.

3.3.2.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group in which to create the consumer.	string
Body	body <i>required</i>	Name and configuration of the consumer. The name is unique within the scope of the consumer group. If a name is not specified, a randomly generated name is assigned. All parameters are optional. The supported configuration options are shown in the following example.	Consumer

3.3.2.3. Responses

HTTP Code	Description	Schema
200	Consumer created successfully.	CreatedConsumer
409	A consumer instance with the specified name already exists in the Kafka Bridge.	Error
422	One or more consumer configuration options have invalid values.	Error

3.3.2.4. Consumes

- **application/vnd.kafka.v2+json**

3.3.2.5. Produces

- **application/vnd.kafka.v2+json**

3.3.2.6. Tags

- Consumers

3.3.2.7. Example HTTP request

3.3.2.7.1. Request body

```
{
  "name" : "consumer1",
  "format" : "binary",
  "auto.offset.reset" : "earliest",
  "enable.auto.commit" : false,
  "fetch.min.bytes" : 512,
  "consumer.request.timeout.ms" : 30000,
  "isolation.level" : "read_committed"
}
```

3.3.2.8. Example HTTP response

3.3.2.8.1. Response 200

```
{
  "instance_id" : "consumer1",
  "base_uri" : "http://localhost:8080/consumers/my-group/instances/consumer1"
}
```

3.3.2.8.2. Response 409

```
{
  "error_code" : 409,
  "message" : "A consumer instance with the specified name already exists in the Kafka Bridge."
}
```

3.3.2.8.3. Response 422

```
{
  "error_code" : 422,
  "message" : "One or more consumer configuration options have invalid values."
}
```

3.3.3. DELETE /consumers/{groupid}/instances/{name}

3.3.3.1. Description

Deletes a specified consumer instance. The request for this operation **MUST** use the base URL (including the host and port) returned in the response from the **POST** request to `/consumers/{groupid}` that was used to create this consumer.

3.3.3.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the consumer belongs.	string
Path	name <i>required</i>	Name of the consumer to delete.	string

3.3.3.3. Responses

HTTP Code	Description	Schema
204	Consumer removed successfully.	No Content
404	The specified consumer instance was not found.	Error

3.3.3.4. Consumes

- **application/vnd.kafka.v2+json**

3.3.3.5. Produces

- **application/vnd.kafka.v2+json**

3.3.3.6. Tags

- Consumers

3.3.3.7. Example HTTP response

3.3.3.7.1. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.4. POST /consumers/{groupid}/instances/{name}/assignments

3.3.4.1. Description

Assigns one or more topic partitions to a consumer.

3.3.4.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the consumer belongs.	string
Path	name <i>required</i>	Name of the consumer to assign topic partitions to.	string
Body	body <i>required</i>	List of topic partitions to assign to the consumer.	Partitions

3.3.4.3. Responses

HTTP Code	Description	Schema
204	Partitions assigned successfully.	No Content
404	The specified consumer instance was not found.	Error
409	Subscriptions to topics, partitions, and patterns are mutually exclusive.	Error

3.3.4.4. Consumes

- **application/vnd.kafka.v2+json**

3.3.4.5. Produces

- **application/vnd.kafka.v2+json**

3.3.4.6. Tags

- Consumers

3.3.4.7. Example HTTP request

3.3.4.7.1. Request body

```
{
  "partitions": [ {
    "topic": "topic",
    "partition": 0
  }, {
    "topic": "topic",
    "partition": 1
  } ]
}
```

3.3.4.8. Example HTTP response

3.3.4.8.1. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.4.8.2. Response 409

```
{
  "error_code" : 409,
  "message" : "Subscriptions to topics, partitions, and patterns are mutually exclusive."
}
```

3.3.5. POST /consumers/{groupid}/instances/{name}/offsets

3.3.5.1. Description

Commits a list of consumer offsets. To commit offsets for all records fetched by the consumer, leave the request body empty.

3.3.5.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the consumer belongs.	string
Path	name <i>required</i>	Name of the consumer.	string
Body	body <i>optional</i>	List of consumer offsets to commit to the consumer offsets commit log. You can specify one or more topic partitions to commit offsets for.	OffsetCommitSeekList

3.3.5.3. Responses

HTTP Code	Description	Schema
204	Commit made successfully.	No Content
404	The specified consumer instance was not found.	Error

3.3.5.4. Consumes

- `application/vnd.kafka.v2+json`

3.3.5.5. Produces

- `application/vnd.kafka.v2+json`

3.3.5.6. Tags

- Consumers

3.3.5.7. Example HTTP request

3.3.5.7.1. Request body

```
{
  "offsets": [ {
    "topic": "topic",
    "partition": 0,
    "offset": 15
  }, {
    "topic": "topic",
    "partition": 1,
    "offset": 42
  } ]
}
```

3.3.5.8. Example HTTP response

3.3.5.8.1. Response 404

```
{
  "error_code": 404,
  "message": "The specified consumer instance was not found."
}
```

3.3.6. POST /consumers/{groupid}/instances/{name}/positions

3.3.6.1. Description

Configures a subscribed consumer to fetch offsets from a particular offset the next time it fetches a set of records from a given topic partition. This overrides the default fetch behavior for consumers. You can specify one or more topic partitions.

3.3.6.2. Parameters

Type	Name	Description	Schema
Path	<code>groupid</code> <i>required</i>	ID of the consumer group to which the consumer belongs.	string

Type	Name	Description	Schema
Path	name <i>required</i>	Name of the subscribed consumer.	string
Body	body <i>required</i>	List of partition offsets from which the subscribed consumer will next fetch records.	OffsetCommitSeekList

3.3.6.3. Responses

HTTP Code	Description	Schema
204	Seek performed successfully.	No Content
404	The specified consumer instance was not found, or the specified consumer instance did not have one of the specified partitions assigned.	Error

3.3.6.4. Consumes

- **application/vnd.kafka.v2+json**

3.3.6.5. Produces

- **application/vnd.kafka.v2+json**

3.3.6.6. Tags

- Consumers
- Seek

3.3.6.7. Example HTTP request

3.3.6.7.1. Request body

```
{
  "offsets": [ {
    "topic": "topic",
    "partition": 0,
    "offset": 15
  }, {
    "topic": "topic",
    "partition": 1,
    "offset": 42
  } ]
}
```


3.3.6.8. Example HTTP response

3.3.6.8.1. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.7. POST /consumers/{groupid}/instances/{name}/positions/beginning

3.3.7.1. Description

Configures a subscribed consumer to seek (and subsequently read from) the first offset in one or more given topic partitions.

3.3.7.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the subscribed consumer belongs.	string
Path	name <i>required</i>	Name of the subscribed consumer.	string
Body	body <i>required</i>	List of topic partitions to which the consumer is subscribed. The consumer will seek the first offset in the specified partitions.	Partitions

3.3.7.3. Responses

HTTP Code	Description	Schema
204	Seek to the beginning performed successfully.	No Content
404	The specified consumer instance was not found, or the specified consumer instance did not have one of the specified partitions assigned.	Error

3.3.7.4. Consumes

- `application/vnd.kafka.v2+json`

3.3.7.5. Produces

- `application/vnd.kafka.v2+json`

3.3.7.6. Tags

- Consumers
- Seek

3.3.7.7. Example HTTP request

3.3.7.7.1. Request body

```
{
  "partitions" : [ {
    "topic" : "topic",
    "partition" : 0
  }, {
    "topic" : "topic",
    "partition" : 1
  } ]
}
```

3.3.7.8. Example HTTP response

3.3.7.8.1. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.8. POST /consumers/{groupid}/instances/{name}/positions/end

3.3.8.1. Description

Configures a subscribed consumer to seek (and subsequently read from) the offset at the end of one or more of the given topic partitions.

3.3.8.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the subscribed consumer belongs.	string
Path	name <i>required</i>	Name of the subscribed consumer.	string
Body	body <i>optional</i>	List of topic partitions to which the consumer is subscribed. The consumer will seek the last offset in the specified partitions.	Partitions

3.3.8.3. Responses

HTTP Code	Description	Schema
204	Seek to the end performed successfully.	No Content
404	The specified consumer instance was not found, or the specified consumer instance did not have one of the specified partitions assigned.	Error

3.3.8.4. Consumes

- `application/vnd.kafka.v2+json`

3.3.8.5. Produces

- `application/vnd.kafka.v2+json`

3.3.8.6. Tags

- Consumers
- Seek

3.3.8.7. Example HTTP request

3.3.8.7.1. Request body

```
{
  "partitions": [ {
    "topic": "topic",
    "partition": 0
  }, {
    "topic": "topic",
    "partition": 1
  } ]
}
```

3.3.8.8. Example HTTP response

3.3.8.8.1. Response 404

```
{
  "error_code": 404,
  "message": "The specified consumer instance was not found."
}
```

3.3.9. GET /consumers/{groupid}/instances/{name}/records

3.3.9.1. Description

Retrieves records for a subscribed consumer, including message values, topics, and partitions. The request for this operation **MUST** use the base URL (including the host and port) returned in the response from the **POST** request to `/consumers/{groupid}` that was used to create this consumer.

3.3.9.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the subscribed consumer belongs.	string
Path	name <i>required</i>	Name of the subscribed consumer to retrieve records from.	string
Query	max_bytes <i>optional</i>	The maximum size, in bytes, of unencoded keys and values that can be included in the response. Otherwise, an error response with code 422 is returned.	integer
Query	timeout <i>optional</i>	The maximum amount of time, in milliseconds, that the HTTP Bridge spends retrieving records before timing out the request.	integer

3.3.9.3. Responses

HTTP Code	Description	Schema
200	Poll request executed successfully.	ConsumerRecordList
404	The specified consumer instance was not found.	Error
406	The format used in the consumer creation request does not match the embedded format in the Accept header of this request or the bridge got a message from the topic which is not JSON encoded.	Error
422	Response exceeds the maximum number of bytes the consumer can receive	Error

3.3.9.4. Produces

- `application/vnd.kafka.json.v2+json`
- `application/vnd.kafka.binary.v2+json`
- `application/vnd.kafka.v2+json`

3.3.9.5. Tags

- Consumers

3.3.9.6. Example HTTP response

3.3.9.6.1. Response 200

```
[ {
  "topic" : "topic",
  "key" : "key1",
  "value" : {
    "foo" : "bar"
  },
  "partition" : 0,
  "offset" : 2
}, {
  "topic" : "topic",
  "key" : "key2",
  "value" : [ "foo2", "bar2" ],
  "partition" : 1,
  "offset" : 3
} ]
```

```
[
  {
    "topic": "test",
    "key": "a2V5",
    "value": "Y29uZmx1ZW50",
    "partition": 1,
    "offset": 100,
  },
  {
    "topic": "test",
    "key": "a2V5",
    "value": "a2Fma2E=",
    "partition": 2,
    "offset": 101,
  }
]
```

3.3.9.6.2. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.9.6.3. Response 406

```
{
  "error_code" : 406,
  "message" : "The `format` used in the consumer creation request does not match the embedded
```

```
format in the Accept header of this request."
```

```
}
```

3.3.9.6.4. Response 422

```
{
  "error_code" : 422,
  "message" : "Response exceeds the maximum number of bytes the consumer can receive"
}
```

3.3.10. POST /consumers/{groupid}/instances/{name}/subscription

3.3.10.1. Description

Subscribes a consumer to one or more topics. You can describe the topics to which the consumer will subscribe in a list (of **Topics** type) or as a **topic_pattern** field. Each call replaces the subscriptions for the subscriber.

3.3.10.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the subscribed consumer belongs.	string
Path	name <i>required</i>	Name of the consumer to subscribe to topics.	string
Body	body <i>required</i>	List of topics to which the consumer will subscribe.	Topics

3.3.10.3. Responses

HTTP Code	Description	Schema
204	Consumer subscribed successfully.	No Content
404	The specified consumer instance was not found.	Error
409	Subscriptions to topics, partitions, and patterns are mutually exclusive.	Error
422	A list (of Topics type) or a topic_pattern must be specified.	Error

3.3.10.4. Consumes

- **application/vnd.kafka.v2+json**

3.3.10.5. Produces

- `application/vnd.kafka.v2+json`

3.3.10.6. Tags

- Consumers

3.3.10.7. Example HTTP request

3.3.10.7.1. Request body

```
{
  "topics": [ "topic1", "topic2" ]
}
```

3.3.10.8. Example HTTP response

3.3.10.8.1. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.10.8.2. Response 409

```
{
  "error_code" : 409,
  "message" : "Subscriptions to topics, partitions, and patterns are mutually exclusive."
}
```

3.3.10.8.3. Response 422

```
{
  "error_code" : 422,
  "message" : "A list (of Topics type) or a topic_pattern must be specified."
}
```

3.3.11. GET /consumers/{groupid}/instances/{name}/subscription

3.3.11.1. Description

Retrieves a list of the topics to which the consumer is subscribed.

3.3.11.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the subscribed consumer belongs.	string
Path	name <i>required</i>	Name of the subscribed consumer.	string

3.3.11.3. Responses

HTTP Code	Description	Schema
200	List of subscribed topics and partitions.	SubscribedTopicList
404	The specified consumer instance was not found.	Error

3.3.11.4. Produces

- **application/vnd.kafka.v2+json**

3.3.11.5. Tags

- Consumers

3.3.11.6. Example HTTP response

3.3.11.6.1. Response 200

```
{
  "topics" : [ "my-topic1", "my-topic2" ],
  "partitions" : [ {
    "my-topic1" : [ 1, 2, 3 ]
  }, {
    "my-topic2" : [ 1 ]
  } ]
}
```

3.3.11.6.2. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```


3.3.12. DELETE /consumers/{groupid}/instances/{name}/subscription

3.3.12.1. Description

Unsubscribes a consumer from all topics.

3.3.12.2. Parameters

Type	Name	Description	Schema
Path	groupid <i>required</i>	ID of the consumer group to which the subscribed consumer belongs.	string
Path	name <i>required</i>	Name of the consumer to unsubscribe from topics.	string

3.3.12.3. Responses

HTTP Code	Description	Schema
204	Consumer unsubscribed successfully.	No Content
404	The specified consumer instance was not found.	Error

3.3.12.4. Tags

- Consumers

3.3.12.5. Example HTTP response

3.3.12.5.1. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified consumer instance was not found."
}
```

3.3.13. GET /healthy

3.3.13.1. Description

Check if the bridge is running. This does not necessarily imply that it is ready to accept requests.

3.3.13.2. Responses

HTTP Code	Description	Schema
200	The bridge is healthy	No Content

3.3.14. GET /openapi

3.3.14.1. Description

Retrieves the OpenAPI v2 specification in JSON format.

3.3.14.2. Responses

HTTP Code	Description	Schema
200	OpenAPI v2 specification in JSON format retrieved successfully.	string

3.3.14.3. Produces

- `application/json`

3.3.15. GET /ready

3.3.15.1. Description

Check if the bridge is ready and can accept requests.

3.3.15.2. Responses

HTTP Code	Description	Schema
200	The bridge is ready	No Content

3.3.16. GET /topics

3.3.16.1. Description

Retrieves a list of all topics.

3.3.16.2. Responses

HTTP Code	Description	Schema
200	List of topics.	< string > array

3.3.16.3. Produces

- `application/vnd.kafka.v2+json`

3.3.16.4. Tags

- Topics

3.3.16.5. Example HTTP response

3.3.16.5.1. Response 200

```
[ "topic1", "topic2" ]
```

3.3.17. POST /topics/{topicname}

3.3.17.1. Description

Sends one or more records to a given topic, optionally specifying a partition, key, or both.

3.3.17.2. Parameters

Type	Name	Description	Schema
Path	topicname <i>required</i>	Name of the topic to send records to or retrieve metadata from.	string
Body	body <i>required</i>		ProducerRecordList

3.3.17.3. Responses

HTTP Code	Description	Schema
200	Records sent successfully.	OffsetRecordSentList
404	The specified topic was not found.	Error
422	The record list is not valid.	Error

3.3.17.4. Consumes

- `application/vnd.kafka.json.v2+json`
- `application/vnd.kafka.binary.v2+json`

3.3.17.5. Produces

- `application/vnd.kafka.v2+json`

3.3.17.6. Tags

- Producer
- Topics

3.3.17.7. Example HTTP request

3.3.17.7.1. Request body

```
{
  "records": [ {
    "key": "key1",
    "value": "value1"
  }, {
    "value": "value2",
    "partition": 1
  }, {
    "value": "value3"
  } ]
}
```

3.3.17.8. Example HTTP response

3.3.17.8.1. Response 200

```
{
  "offsets": [ {
    "partition": 2,
    "offset": 0
  }, {
    "partition": 1,
    "offset": 1
  }, {
    "partition": 2,
    "offset": 2
  } ]
}
```

3.3.17.8.2. Response 404

```
{
```

```

    "error_code" : 404,
    "message" : "The specified topic was not found."
  }

```

3.3.17.8.3. Response 422

```

{
  "error_code" : 422,
  "message" : "The record list contains invalid records."
}

```

3.3.18. GET /topics/{topicname}

3.3.18.1. Description

Retrieves the metadata about a given topic.

3.3.18.2. Parameters

Type	Name	Description	Schema
Path	topicname <i>required</i>	Name of the topic to send records to or retrieve metadata from.	string

3.3.18.3. Responses

HTTP Code	Description	Schema
200	Topic metadata	TopicMetadata

3.3.18.4. Produces

- **application/vnd.kafka.v2+json**

3.3.18.5. Tags

- Topics

3.3.18.6. Example HTTP response

3.3.18.6.1. Response 200

```

{
  "name" : "topic",
  "offset" : 2,
  "configs" : {
    "cleanup.policy" : "compact"
  }
}

```

```

    },
    "partitions" : [ {
      "partition" : 1,
      "leader" : 1,
      "replicas" : [ {
        "broker" : 1,
        "leader" : true,
        "in_sync" : true
      }, {
        "broker" : 2,
        "leader" : false,
        "in_sync" : true
      } ]
    }, {
      "partition" : 2,
      "leader" : 2,
      "replicas" : [ {
        "broker" : 1,
        "leader" : false,
        "in_sync" : true
      }, {
        "broker" : 2,
        "leader" : true,
        "in_sync" : true
      } ]
    } ]
  } ]
}

```

3.3.19. GET /topics/{topicname}/partitions

3.3.19.1. Description

Retrieves a list of partitions for the topic.

3.3.19.2. Parameters

Type	Name	Description	Schema
Path	topicname <i>required</i>	Name of the topic to send records to or retrieve metadata from.	string

3.3.19.3. Responses

HTTP Code	Description	Schema
200	List of partitions	< PartitionMetadata > array
404	The specified topic was not found.	Error

3.3.19.4. Produces

- `application/vnd.kafka.v2+json`

3.3.19.5. Tags

- Topics

3.3.19.6. Example HTTP response

3.3.19.6.1. Response 200

```
[ {
  "partition" : 1,
  "leader" : 1,
  "replicas" : [ {
    "broker" : 1,
    "leader" : true,
    "in_sync" : true
  }, {
    "broker" : 2,
    "leader" : false,
    "in_sync" : true
  } ]
}, {
  "partition" : 2,
  "leader" : 2,
  "replicas" : [ {
    "broker" : 1,
    "leader" : false,
    "in_sync" : true
  }, {
    "broker" : 2,
    "leader" : true,
    "in_sync" : true
  } ]
} ]
```

3.3.19.6.2. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified topic was not found."
}
```

3.3.20. POST /topics/{topicname}/partitions/{partitionid}

3.3.20.1. Description

Sends one or more records to a given topic partition, optionally specifying a key.

3.3.20.2. Parameters

Type	Name	Description	Schema
Path	partitionid <i>required</i>	ID of the partition to send records to or retrieve metadata from.	integer
Path	topicname <i>required</i>	Name of the topic to send records to or retrieve metadata from.	string
Body	body <i>required</i>	List of records to send to a given topic partition, including a value (required) and a key (optional).	ProducerRecordToPartitionList

3.3.20.3. Responses

HTTP Code	Description	Schema
200	Records sent successfully.	OffsetRecordSentList
404	The specified topic partition was not found.	Error
422	The record is not valid.	Error

3.3.20.4. Consumes

- **application/vnd.kafka.json.v2+json**
- **application/vnd.kafka.binary.v2+json**

3.3.20.5. Produces

- **application/vnd.kafka.v2+json**

3.3.20.6. Tags

- Producer
- Topics

3.3.20.7. Example HTTP request

3.3.20.7.1. Request body

```
{
  "records": [ {
    "key": "key1",
    "value": "value1"
  }, {
```



```

    "value" : "value2"
  }
}

```

3.3.20.8. Example HTTP response

3.3.20.8.1. Response 200

```

{
  "offsets" : [ {
    "partition" : 2,
    "offset" : 0
  }, {
    "partition" : 1,
    "offset" : 1
  }, {
    "partition" : 2,
    "offset" : 2
  }
]
}

```

3.3.20.8.2. Response 404

```

{
  "error_code" : 404,
  "message" : "The specified topic partition was not found."
}

```

3.3.20.8.3. Response 422

```

{
  "error_code" : 422,
  "message" : "The record is not valid."
}

```

3.3.21. GET /topics/{topicname}/partitions/{partitionid}

3.3.21.1. Description

Retrieves partition metadata for the topic partition.

3.3.21.2. Parameters

Type	Name	Description	Schema
Path	partitionid <i>required</i>	ID of the partition to send records to or retrieve metadata from.	integer

Type	Name	Description	Schema
Path	topicname <i>required</i>	Name of the topic to send records to or retrieve metadata from.	string

3.3.21.3. Responses

HTTP Code	Description	Schema
200	Partition metadata	PartitionMetadata
404	The specified topic partition was not found.	Error

3.3.21.4. Produces

- **application/vnd.kafka.v2+json**

3.3.21.5. Tags

- Topics

3.3.21.6. Example HTTP response

3.3.21.6.1. Response 200

```
{
  "partition" : 1,
  "leader" : 1,
  "replicas" : [ {
    "broker" : 1,
    "leader" : true,
    "in_sync" : true
  }, {
    "broker" : 2,
    "leader" : false,
    "in_sync" : true
  } ]
}
```

3.3.21.6.2. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified topic partition was not found."
}
```

3.3.22. GET /topics/{topicname}/partitions/{partitionid}/offsets

3.3.22.1. Description

Retrieves a summary of the offsets for the topic partition.

3.3.22.2. Parameters

Type	Name	Description	Schema
Path	partitionid <i>required</i>	ID of the partition.	integer
Path	topicname <i>required</i>	Name of the topic containing the partition.	string

3.3.22.3. Responses

HTTP Code	Description	Schema
200	A summary of the offsets for the topic partition.	OffsetsSummary
404	The specified topic partition was not found.	Error

3.3.22.4. Produces

- **application/vnd.kafka.v2+json**

3.3.22.5. Tags

- Topics

3.3.22.6. Example HTTP response

3.3.22.6.1. Response 200

```
{
  "beginning_offset" : 10,
  "end_offset" : 50
}
```

3.3.22.6.2. Response 404

```
{
  "error_code" : 404,
  "message" : "The specified topic partition was not found."
}
```

APPENDIX A. USING YOUR SUBSCRIPTION

AMQ Streams is provided through a software subscription. To manage your subscriptions, access your account at the Red Hat Customer Portal.

Accessing Your Account

1. Go to access.redhat.com.
2. If you do not already have an account, create one.
3. Log in to your account.

Activating a Subscription

1. Go to access.redhat.com.
2. Navigate to **My Subscriptions**.
3. Navigate to **Activate a subscription** and enter your 16-digit activation number.

Downloading Zip and Tar Files

To access zip or tar files, use the customer portal to find the relevant files for download. If you are using RPM packages, this step is not required.

1. Open a browser and log in to the Red Hat Customer Portal **Product Downloads** page at access.redhat.com/downloads.
2. Locate the **AMQ Streams for Apache Kafka** entries in the **INTEGRATION AND AUTOMATION** category.
3. Select the desired AMQ Streams product. The **Software Downloads** page opens.
4. Click the **Download** link for your component.

Revised on 2022-04-13 17:50:11 UTC