



Red Hat AMQ 2020.Q4

AMQ Streams on OpenShift Overview

For use with AMQ Streams 1.6 on OpenShift Container Platform

Red Hat AMQ 2020.Q4 AMQ Streams on OpenShift Overview

For use with AMQ Streams 1.6 on OpenShift Container Platform

Legal Notice

Copyright © 2020 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

This guide provides an overview of the features and capabilities of AMQ Streams.

Table of Contents

CHAPTER 1. KEY FEATURES	4
1.1. KAFKA CAPABILITIES	4
1.2. KAFKA USE CASES	4
1.3. HOW AMQ STREAMS SUPPORTS KAFKA	4
CHAPTER 2. ABOUT KAFKA	6
2.1. KAFKA CONCEPTS	6
2.2. PRODUCERS AND CONSUMERS	7
CHAPTER 3. AMQ STREAMS DEPLOYMENT OF KAFKA	9
3.1. KAFKA COMPONENT ARCHITECTURE	9
3.2. KAFKA BRIDGE INTERFACE	10
3.2.1. HTTP requests	10
3.2.2. Supported clients for the Kafka Bridge	11
CHAPTER 4. AMQ STREAMS OPERATORS	13
Operators	13
4.1. CLUSTER OPERATOR	14
4.2. TOPIC OPERATOR	15
4.3. USER OPERATOR	16
CHAPTER 5. KAFKA CONFIGURATION	17
5.1. CUSTOM RESOURCES	17
Kafka topic custom resource	17
5.2. COMMON CONFIGURATION	17
Example YAML showing common configuration	18
5.3. KAFKA CLUSTER CONFIGURATION	19
Example YAML showing Kafka configuration	20
5.4. KAFKA MIRRORMAKER CONFIGURATION	20
MirrorMaker 2.0	20
Cluster configuration	21
Bidirectional replication across two clusters	22
Example YAML showing MirrorMaker 2.0 configuration	22
MirrorMaker	23
Key Consumer configuration	23
Key Producer configuration	23
Example YAML showing MirrorMaker configuration	23
5.5. KAFKA CONNECT CONFIGURATION	24
Example YAML showing Kafka Connect configuration	24
Connectors	24
Managing connectors	25
Example YAML showing KafkaConnector configuration	25
Example YAML showing annotation to enable KafkaConnector	25
5.6. KAFKA BRIDGE CONFIGURATION	25
CORS	26
Example YAML showing Kafka Bridge configuration	26
CHAPTER 6. SECURING KAFKA	27
6.1. ENCRYPTION	27
6.2. AUTHENTICATION	27
6.3. AUTHORIZATION	28
CHAPTER 7. MONITORING	29

7.1. PROMETHEUS	29
7.2. GRAFANA	30
7.3. KAFKA EXPORTER	30
7.4. DISTRIBUTED TRACING	30
Tracing for Kafka clients	30
7.5. CRUISE CONTROL	30
APPENDIX A. USING YOUR SUBSCRIPTION	31
Accessing Your Account	31
Activating a Subscription	31
Downloading Zip and Tar Files	31

CHAPTER 1. KEY FEATURES

AMQ Streams simplifies the process of running Apache Kafka in an OpenShift cluster.

This guide is intended as a starting point for building an understanding of AMQ Streams. The guide introduces some of the key concepts behind Kafka, which is central to AMQ Streams, explaining briefly the purpose of Kafka components. Configuration points are outlined, including options to secure and monitor Kafka. A distribution of AMQ Streams provides the files to deploy and manage a Kafka cluster, as well as example files for configuration and monitoring of your deployment.

A typical Kafka deployment is described, as well as the tools used to deploy and manage Kafka.

1.1. KAFKA CAPABILITIES

The underlying data stream-processing capabilities and component architecture of Kafka can deliver:

- Microservices and other applications to share data with extremely high throughput and low latency
- Message ordering guarantees
- Message rewind/replay from data storage to reconstruct an application state
- Message compaction to remove old records when using a key-value log
- Horizontal scalability in a cluster configuration
- Replication of data to control fault tolerance
- Retention of high volumes of data for immediate access

1.2. KAFKA USE CASES

Kafka's capabilities make it suitable for:

- Event-driven architectures
- Event sourcing to capture changes to the state of an application as a log of events
- Message brokering
- Website activity tracking
- Operational monitoring through metrics
- Log collection and aggregation
- Commit logs for distributed systems
- Stream processing so that applications can respond to data in real time

1.3. HOW AMQ STREAMS SUPPORTS KAFKA

AMQ Streams provides container images and Operators for running Kafka on OpenShift. AMQ Streams Operators are fundamental to the running of AMQ Streams. The Operators provided with AMQ Streams are purpose-built with specialist operational knowledge to effectively manage Kafka.

Operators simplify the process of:

- Deploying and running Kafka clusters
- Deploying and running Kafka components
- Configuring access to Kafka
- Securing access to Kafka
- Upgrading Kafka
- Managing brokers
- Creating and managing topics
- Creating and managing users

CHAPTER 2. ABOUT KAFKA

Apache Kafka is an open-source distributed publish-subscribe messaging system for fault-tolerant real-time data feeds.

Additional resources

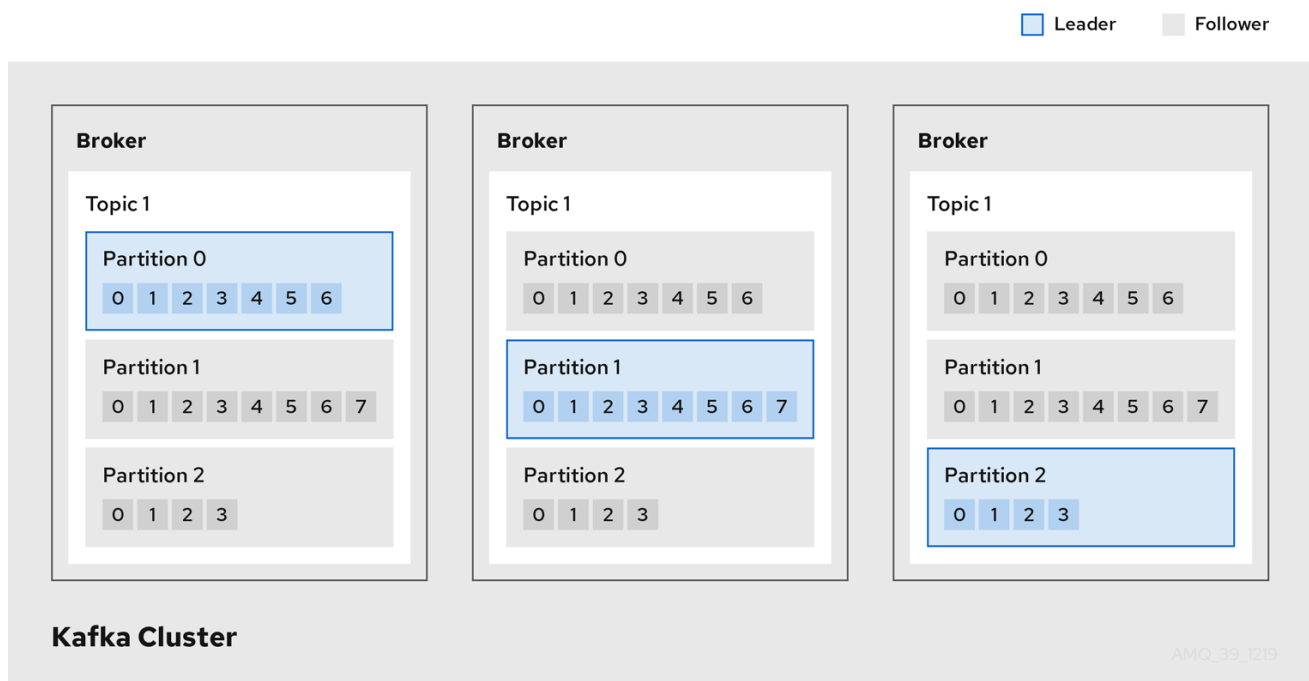
- For more information about Apache Kafka, see the [Apache Kafka website](#).

2.1. KAFKA CONCEPTS

Knowledge of the key concepts of Kafka is important in understanding how AMQ Streams works.

A Kafka cluster comprises multiple brokers. Topics are used to receive and store data in a Kafka cluster. Topics are split by partitions, where the data is written. Partitions are replicated across topics for fault tolerance.

Kafka brokers and topics



Broker

A broker, sometimes referred to as a server or node, orchestrates the storage and passing of messages.

Topic

A topic provides a destination for the storage of data. Each topic is split into one or more partitions.

Cluster

A group of broker instances.

Partition

The number of topic partitions is defined by a topic *partition count*.

Partition leader

A partition leader handles all producer requests for a topic.

Partition follower

A partition follower replicates the partition data of a partition leader, optionally handling consumer requests.

Topics use a *replication factor* to configure the number of replicas of each partition within the cluster. A topic comprises at least one partition.


An *in-sync* replica has the same number of messages as the leader. Configuration defines how many replicas must be *in-sync* to be able to produce messages, ensuring that a message is committed only after it has been successfully copied to the replica partition. In this way, if the leader fails the message is not lost.

In the *Kafka brokers and topics* diagram, we can see each numbered partition has a leader and two followers in replicated topics.

2.2. PRODUCERS AND CONSUMERS

Producers and consumers send and receive messages (publish and subscribe) through brokers. Messages comprise an optional *key* and a *value* that contains the message data, plus headers and related metadata. The key is used to identify the subject of the message, or a property of the message. Messages are delivered in batches, and batches and records contain headers and metadata that provide details that are useful for filtering and routing by clients, such as the timestamp and offset position for the record.

Producers and consumers



AMQ_39_1219

Producer

A producer sends messages to a broker topic to be written to the end offset of a partition. Messages are written to partitions by a producer on a round robin basis, or to a specific partition based on the message key.

Consumer

A consumer subscribes to a topic and reads messages according to topic, partition and offset.

Consumer group

Consumer groups are used to share a typically large data stream generated by multiple producers


from a given topic. Consumers are grouped using a **group.id**, allowing messages to be spread across the members. Consumers within a group do not read data from the same partition, but can receive data from one or more partitions.

Offsets

Offsets describe the position of messages within a partition. Each message in a given partition has a unique offset, which helps identify the position of a consumer within the partition to track the number of records that have been consumed.

Committed offsets are written to an offset commit log. A **__consumer_offsets** topic stores information on committed offsets, the position of last and next offset, according to consumer group.

Producing and consuming data



CHAPTER 3. AMQ STREAMS DEPLOYMENT OF KAFKA

Apache Kafka components are provided for deployment to OpenShift with the AMQ Streams distribution. The Kafka components are generally run as clusters for availability.

A typical deployment incorporating Kafka components might include:

- **Kafka** cluster of broker nodes
- **ZooKeeper** cluster of replicated ZooKeeper instances
- **Kafka Connect** cluster for external data connections
- **Kafka MirrorMaker** cluster to mirror the Kafka cluster in a secondary cluster
- **Kafka Exporter** to extract additional Kafka metrics data for monitoring
- **Kafka Bridge** to make HTTP-based requests to the Kafka cluster

Not all of these components are mandatory, though you need Kafka and ZooKeeper as a minimum. Some components can be deployed without Kafka, such as MirrorMaker or Kafka Connect.


3.1. KAFKA COMPONENT ARCHITECTURE

A cluster of Kafka brokers is main part of the Apache Kafka project responsible for delivering messages.

A broker uses Apache ZooKeeper for storing configuration data and for cluster coordination. Before running Apache Kafka, an Apache ZooKeeper cluster has to be ready.

Each of the other Kafka components interact with the Kafka cluster to perform specific roles.

Kafka component interaction



AMQ_39_1219

Apache ZooKeeper is a core dependency for Kafka as it provides a cluster coordination service, storing and tracking the status of brokers and consumers. ZooKeeper is also used for leader election of partitions.

Kafka Connect

Kafka Connect is an integration toolkit for streaming data between Kafka brokers and other systems using *Connector* plugins. Kafka Connect provides a framework for integrating Kafka with an external data source or target, such as a database, for import or export of data using connectors. Connectors are plugins that provide the connection configuration needed.

- A *source* connector pushes external data into Kafka.
- A *sink* connector extracts data out of Kafka
External data is translated and transformed into the appropriate format.

You can deploy Kafka Connect with Source2Image support, which provides a convenient way to include connectors.

Kafka MirrorMaker

Kafka MirrorMaker replicates data between two Kafka clusters, within or across data centers. MirrorMaker takes messages from a source Kafka cluster and writes them to a target Kafka cluster.

Kafka Bridge

Kafka Bridge provides an API for integrating HTTP-based clients with a Kafka cluster.

Kafka Exporter

Kafka Exporter extracts data for analysis as Prometheus metrics, primarily data relating to offsets, consumer groups, consumer lag and topics. Consumer lag is the delay between the last message written to a partition and the message currently being picked up from that partition by a consumer

3.2. 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.

3.2.1. 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

- To view the API documentation, including example requests and responses, see the [Kafka Bridge API reference](#).

3.2.2. Supported clients for the Kafka Bridge

You can use the Kafka Bridge to integrate both *internal* and *external* HTTP client applications with your Kafka cluster.

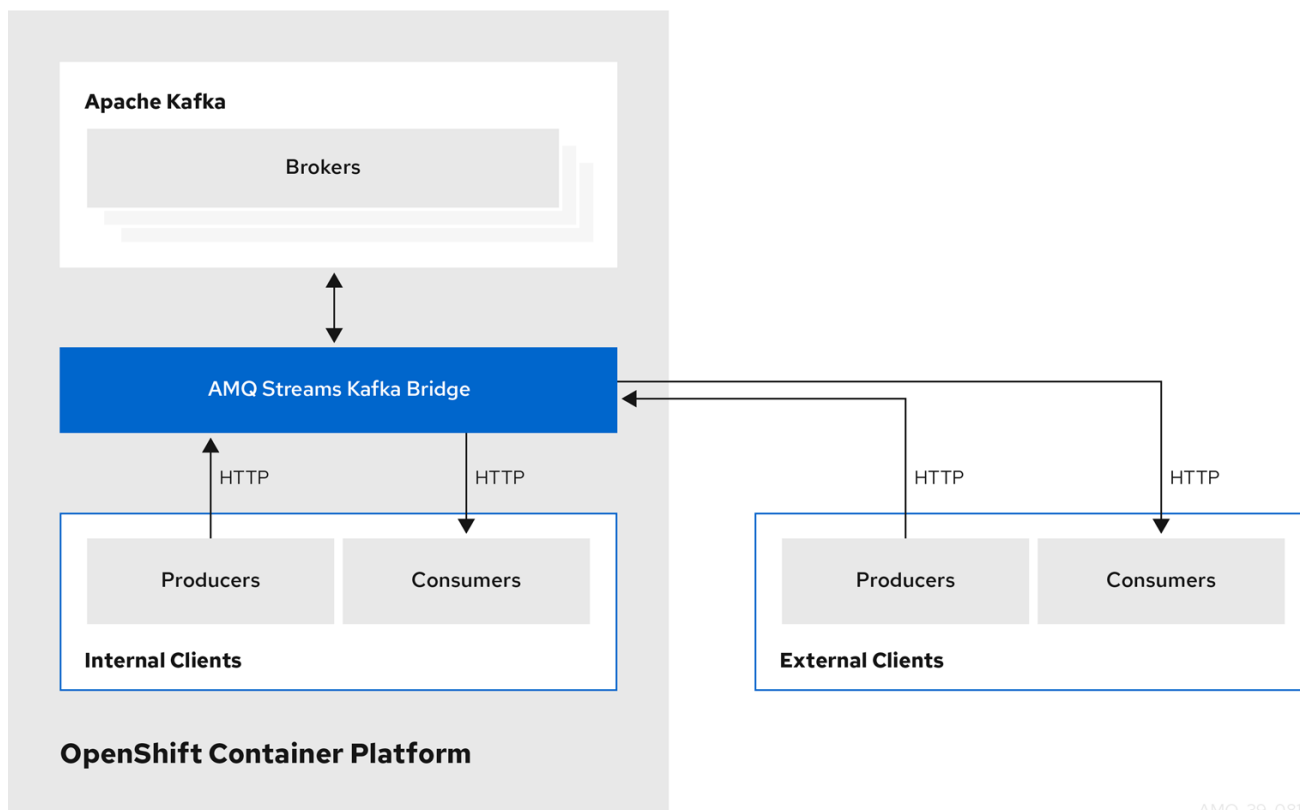
Internal clients

Internal clients are container-based HTTP clients running in *the same* OpenShift cluster as the Kafka Bridge itself. Internal clients can access the Kafka Bridge on the host and port defined in the **KafkaBridge** custom resource.

External clients

External clients are HTTP clients running *outside* the OpenShift cluster in which the Kafka Bridge is deployed and running. External clients can access the Kafka Bridge through an OpenShift Route, a loadbalancer service, or using an Ingress.

HTTP internal and external client integration



AMQ_39_0819

CHAPTER 4. AMQ STREAMS OPERATORS

AMQ Streams supports Kafka using *Operators* to deploy and manage the components and dependencies of Kafka to OpenShift.

Operators are a method of packaging, deploying, and managing an OpenShift application. AMQ Streams Operators extend OpenShift functionality, automating common and complex tasks related to a Kafka deployment. By implementing knowledge of Kafka operations in code, Kafka administration tasks are simplified and require less manual intervention.

Operators

AMQ Streams provides Operators for managing a Kafka cluster running within an OpenShift cluster.

Cluster Operator

Deploys and manages Apache Kafka clusters, Kafka Connect, Kafka MirrorMaker, Kafka Bridge, Kafka Exporter, and the Entity Operator

Entity Operator

Comprises the Topic Operator and User Operator

Topic Operator

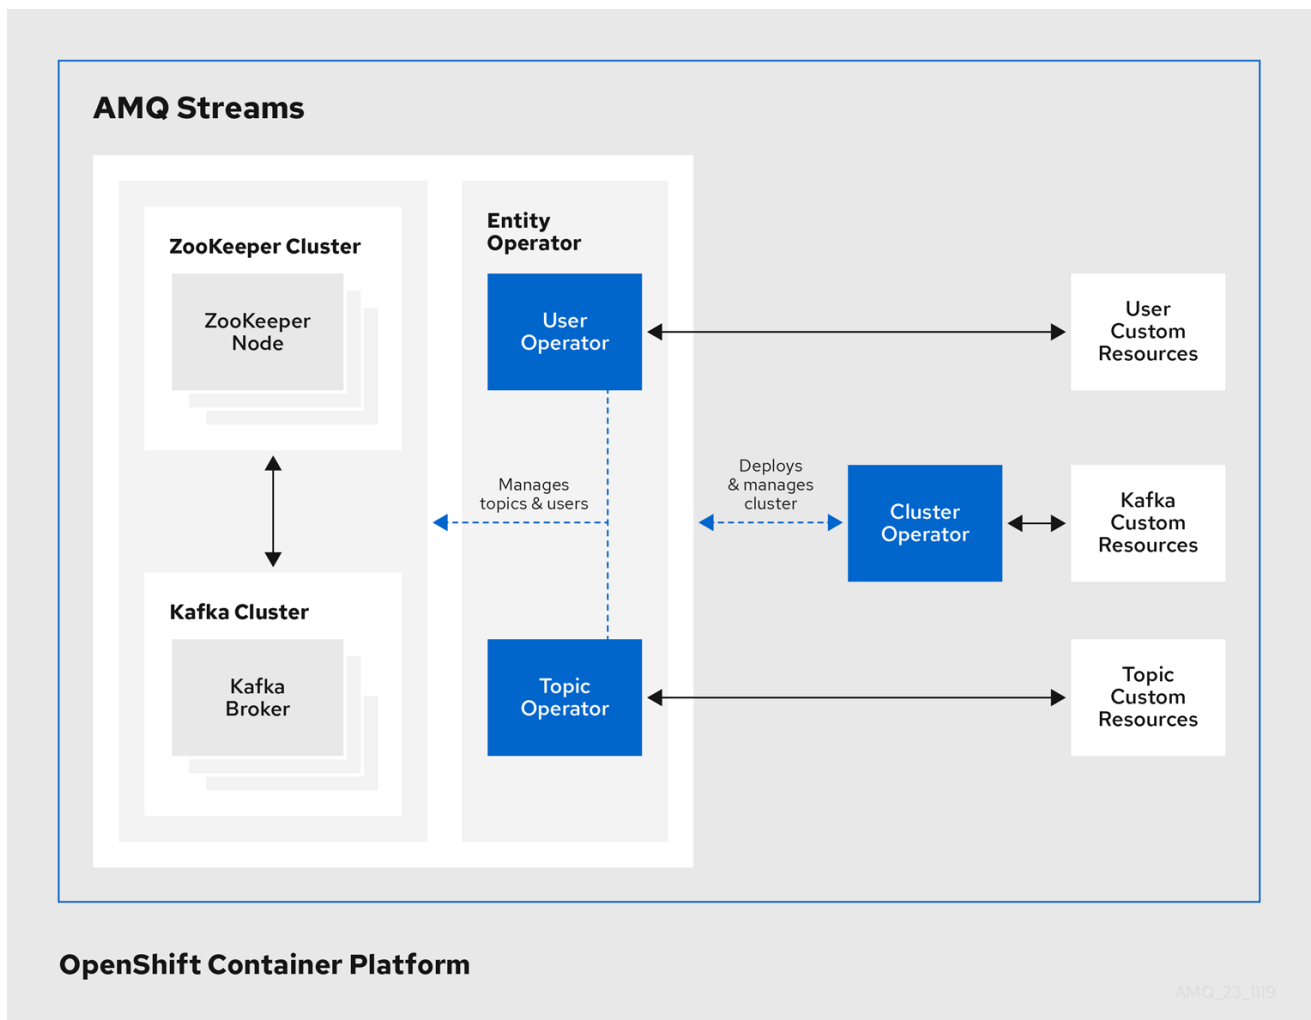
Manages Kafka topics

User Operator

Manages Kafka users

The Cluster Operator can deploy the Topic Operator and User Operator as part of an **Entity Operator** configuration at the same time as a Kafka cluster.

Operators within the AMQ Streams architecture



4.1. CLUSTER OPERATOR

AMQ Streams uses the Cluster Operator to deploy and manage clusters for:

- Kafka (including ZooKeeper, Entity Operator, Kafka Exporter, and Cruise Control)
- Kafka Connect
- Kafka MirrorMaker
- Kafka Bridge

Custom resources are used to deploy the clusters.

For example, to deploy a Kafka cluster:

- A **Kafka** resource with the cluster configuration is created within the OpenShift cluster.
- The Cluster Operator deploys a corresponding Kafka cluster, based on what is declared in the **Kafka** resource.


The Cluster Operator can also deploy (through configuration of the **Kafka** resource):

- A Topic Operator to provide operator-style topic management through **KafkaTopic** custom resources

- A User Operator to provide operator-style user management through **KafkaUser** custom resources

The Topic Operator and User Operator function within the Entity Operator on deployment.

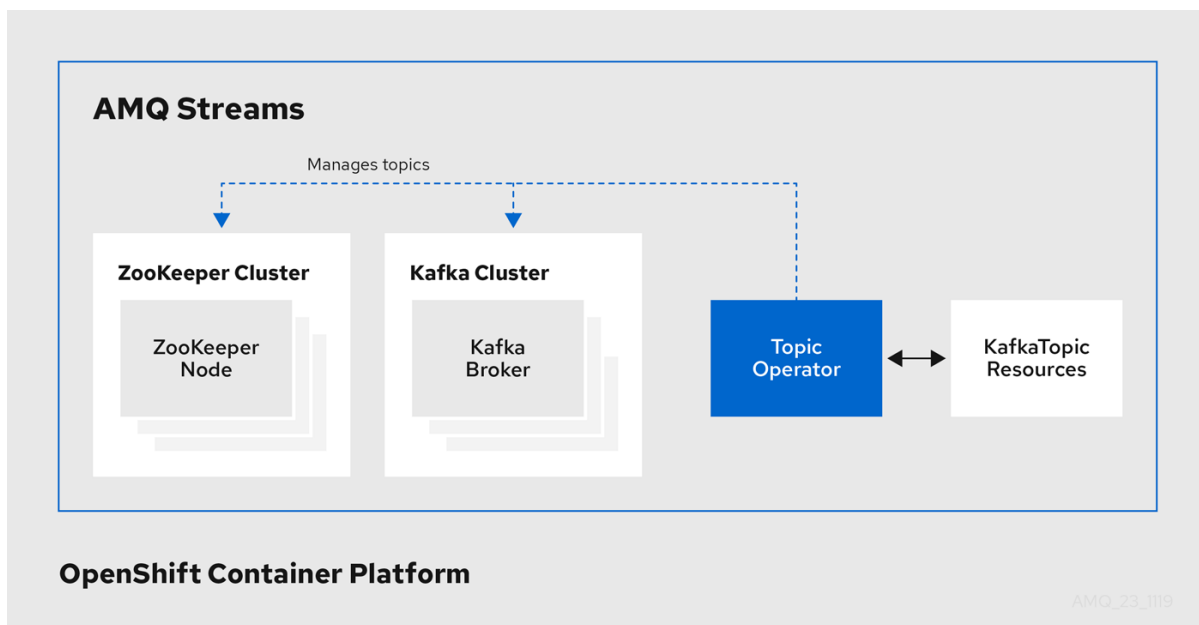
Example architecture for the Cluster Operator



4.2. TOPIC OPERATOR

The Topic Operator provides a way of managing topics in a Kafka cluster through OpenShift resources.

Example architecture for the Topic Operator



The role of the Topic Operator is to keep a set of **KafkaTopic** OpenShift resources describing Kafka topics in-sync with corresponding Kafka topics.

Specifically, if a **KafkaTopic** is:

- Created, the Topic Operator creates the topic

- Deleted, the Topic Operator deletes the topic
- Changed, the Topic Operator updates the topic

Working in the other direction, if a topic is:

- Created within the Kafka cluster, the Operator creates a **KafkaTopic**
- Deleted from the Kafka cluster, the Operator deletes the **KafkaTopic**
- Changed in the Kafka cluster, the Operator updates the **KafkaTopic**

This allows you to declare a **KafkaTopic** as part of your application's deployment and the Topic Operator will take care of creating the topic for you. Your application just needs to deal with producing or consuming from the necessary topics.

If the topic is reconfigured or reassigned to different Kafka nodes, the **KafkaTopic** will always be up to date.

4.3. USER OPERATOR

The User Operator manages Kafka users for a Kafka cluster by watching for **KafkaUser** resources that describe Kafka users, and ensuring that they are configured properly in the Kafka cluster.

For example, if a **KafkaUser** is:

- Created, the User Operator creates the user it describes
- Deleted, the User Operator deletes the user it describes
- Changed, the User Operator updates the user it describes

Unlike the Topic Operator, the User Operator does not sync any changes from the Kafka cluster with the OpenShift resources. Kafka topics can be created by applications directly in Kafka, but it is not expected that the users will be managed directly in the Kafka cluster in parallel with the User Operator.

The User Operator allows you to declare a **KafkaUser** resource as part of your application's deployment. You can specify the authentication and authorization mechanism for the user. You can also configure *user quotas* that control usage of Kafka resources to ensure, for example, that a user does not monopolize access to a broker.

When the user is created, the user credentials are created in a **Secret**. Your application needs to use the user and its credentials for authentication and to produce or consume messages.

In addition to managing credentials for authentication, the User Operator also manages authorization rules by including a description of the user's access rights in the **KafkaUser** declaration.

CHAPTER 5. KAFKA CONFIGURATION

A deployment of Kafka components to an OpenShift cluster using AMQ Streams is highly configurable through the application of custom resources. Custom resources are created as instances of APIs added by Custom resource definitions (CRDs) to extend OpenShift resources.

CRDs act as configuration instructions to describe the custom resources in an OpenShift cluster, and are provided with AMQ Streams for each Kafka component used in a deployment, as well as users and topics. CRDs and custom resources are defined as YAML files. Example YAML files are provided with the AMQ Streams distribution.

CRDs also allow AMQ Streams resources to benefit from native OpenShift features like CLI accessibility and configuration validation.

In this chapter we look at how Kafka components are configured through custom resources, starting with common configuration points and then important configuration considerations specific to components.

5.1. CUSTOM RESOURCES

After a new custom resource type is added to your cluster by installing a CRD, you can create instances of the resource based on its specification.

The custom resources for AMQ Streams components have common configuration properties, which are defined under **spec**.

In this fragment from a Kafka topic custom resource, the **apiVersion** and **kind** properties identify the associated CRD. The **spec** property shows configuration that defines the number of partitions and replicas for the topic.

Kafka topic custom resource

```
apiVersion: kafka.strimzi.io/v1beta1
kind: KafkaTopic
metadata:
  name: my-topic
  labels:
    strimzi.io/cluster: my-cluster
spec:
  partitions: 1
  replicas: 1
# ...
```

There are many additional configuration options that can be incorporated into a YAML definition, some common and some specific to a particular component.

5.2. COMMON CONFIGURATION

Some of the configuration options common to resources are described here. [Security](#) and [metrics collection](#) might also be adopted where applicable.

Bootstrap servers

Bootstrap servers are used for host/port connection to a Kafka cluster for:

- Kafka Connect

- Kafka Bridge
- Kafka MirrorMaker producers and consumers

CPU and memory resources

You request CPU and memory resources for components. Limits specify the maximum resources that can be consumed by a given container.

Resource requests and limits for the Topic Operator and User Operator are set in the **Kafka** resource.

Logging

You define the logging level for the component. Logging can be defined directly (inline) or externally using a config map.

Healthchecks

Healthcheck configuration introduces *liveness* and *readiness* probes to know when to restart a container (liveness) and when a container can accept traffic (readiness).

JVM options

JVM options provide maximum and minimum memory allocation to optimize the performance of the component according to the platform it is running on.

Pod scheduling

Pod schedules use *affinity/anti-affinity* rules to determine under what circumstances a pod is scheduled onto a node.

Example YAML showing common configuration

```
apiVersion: kafka.strimzi.io/v1beta1
kind: KafkaConnect
metadata:
  name: my-cluster
spec:
  # ...
  bootstrapServers: my-cluster-kafka-bootstrap:9092
  resources:
    requests:
      cpu: 12
      memory: 64Gi
    limits:
      cpu: 12
      memory: 64Gi
  logging:
    type: inline
    loggers:
      connect.root.logger.level: "INFO"
  readinessProbe:
    initialDelaySeconds: 15
    timeoutSeconds: 5
  livenessProbe:
    initialDelaySeconds: 15
    timeoutSeconds: 5
  jvmOptions:
    "-Xmx": "2g"
    "-Xms": "2g"
  template:
```

```

pod:
  affinity:
    nodeAffinity:
      requiredDuringSchedulingIgnoredDuringExecution:
        nodeSelectorTerms:
          - matchExpressions:
              - key: node-type
                operator: In
              values:
                - fast-network
# ...

```

5.3. KAFKA CLUSTER CONFIGURATION

A kafka cluster comprises one or more brokers. For producers and consumers to be able to access topics within the brokers, Kafka configuration must define how data is stored in the cluster, and how the data is accessed. You can configure a Kafka cluster to run with multiple broker nodes across *racks*.

Storage

Kafka and ZooKeeper store data on disks.

AMQ Streams requires block storage provisioned through **StorageClass**. The file system format for storage must be *XFS* or *EXT4*. Three types of data storage are supported:

Ephemeral (Recommended for development only)

Ephemeral storage stores data for the lifetime of an instance. Data is lost when the instance is restarted.

Persistent

Persistent storage relates to long-term data storage independent of the lifecycle of the instance.

JBOD (Just a Bunch of Disks, suitable for Kafka only)

JBOD allows you to use multiple disks to store commit logs in each broker.

The disk capacity used by an existing Kafka cluster can be increased if supported by the infrastructure.

Listeners

Listeners configure how clients connect to a Kafka cluster.

By specifying a unique name and port for each listener within a Kafka cluster, you can configure multiple listeners.

The following types of listener are supported:

- **Internal listeners** for access within OpenShift
- **External listeners** for access outside of OpenShift

You can enable TLS encryption for listeners, and configure [authentication](#).

Internal listeners are specified using an **internal** type.

External listeners expose Kafka by specifying an external **type**:

- **route** to use OpenShift routes and the default HAProxy router

- **loadbalancer** to use loadbalancer services
- **nodeport** to use ports on OpenShift nodes
- **ingress** to use OpenShift *Ingress* and the [NGINX Ingress Controller for Kubernetes](#)

If you are using [OAuth 2.0 for token-based authentication](#), you can configure listeners to use the authorization server.

Rack awareness

Rack awareness is a configuration feature that distributes Kafka broker pods and topic replicas across *racks*, which represent data centers or racks in data centers, or availability zones.

Example YAML showing Kafka configuration

```
apiVersion: kafka.strimzi.io/v1beta1
kind: Kafka
metadata:
  name: my-cluster
spec:
  kafka:
    # ...
    listeners:
      - name: tls
        port: 9093
        type: internal
        tls: true
        authentication:
          type: tls
      - name: external1
        port: 9094
        type: route
        tls: true
        authentication:
          type: tls
    # ...
    storage:
      type: persistent-claim
      size: 10000Gi
    # ...
    rack:
      topologyKey: topology.kubernetes.io/zone
    # ...
```

5.4. KAFKA MIRRORMAKER CONFIGURATION

To set up MirrorMaker, a source and target (destination) Kafka cluster must be running.

You can use AMQ Streams with MirrorMaker 2.0, although the earlier version of MirrorMaker continues to be supported.

MirrorMaker 2.0

MirrorMaker 2.0 is based on the Kafka Connect framework, *connectors* managing the transfer of data between clusters.

MirrorMaker 2.0 uses:

- Source cluster configuration to consume data from the source cluster
- Target cluster configuration to output data to the target cluster

Cluster configuration

You can use MirrorMaker 2.0 in *active/passive* or *active/active* cluster configurations.

- In an *active/active* configuration, both clusters are active and provide the same data simultaneously, which is useful if you want to make the same data available locally in different geographical locations.
- In an *active/passive* configuration, the data from an active cluster is replicated in a passive cluster, which remains on standby, for example, for data recovery in the event of system failure.

You configure a **KafkaMirrorMaker2** custom resource to define the Kafka Connect deployment, including the connection details of the source and target clusters, and then run a set of MirrorMaker 2.0 connectors to make the connection.

Topic configuration is automatically synchronized between the source and target clusters according to the topics defined in the **KafkaMirrorMaker2** custom resource. Configuration changes are propagated to remote topics so that new topics and partitions are detected and created. Topic replication is defined using regular expression patterns to *whitelist* or *blacklist* topics.

The following MirrorMaker 2.0 connectors and related internal topics help manage the transfer and synchronization of data between the clusters.

MirrorSourceConnector

A *MirrorSourceConnector* creates remote topics from the source cluster.

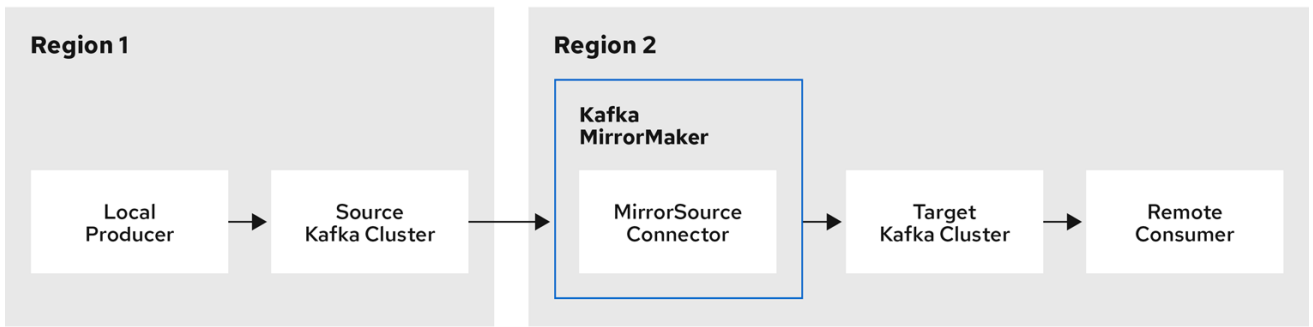
MirrorCheckpointConnector

A *MirrorCheckpointConnector* tracks and maps offsets for specified consumer groups using an *offset sync* topic and *checkpoint* topic. The offset sync topic maps the source and target offsets for replicated topic partitions from record metadata. A checkpoint is emitted from each source cluster and replicated in the target cluster through the checkpoint topic. The checkpoint topic maps the last committed offset in the source and target cluster for replicated topic partitions in each consumer group.

MirrorHeartbeatConnector

A *MirrorHeartbeatConnector* periodically checks connectivity between clusters. A heartbeat is produced every second by the *MirrorHeartbeatConnector* into a *heartbeat* topic that is created on the local cluster. If you have MirrorMaker 2.0 at both the remote and local locations, the heartbeat emitted at the remote location by the *MirrorHeartbeatConnector* is treated like any remote topic and mirrored by the *MirrorSourceConnector* at the local cluster. The heartbeat topic makes it easy to check that the remote cluster is available and the clusters are connected. If things go wrong, the heartbeat topic offset positions and time stamps can help with recovery and diagnosis.

Figure 5.1. Replication across two clusters



AMQ_73_0220

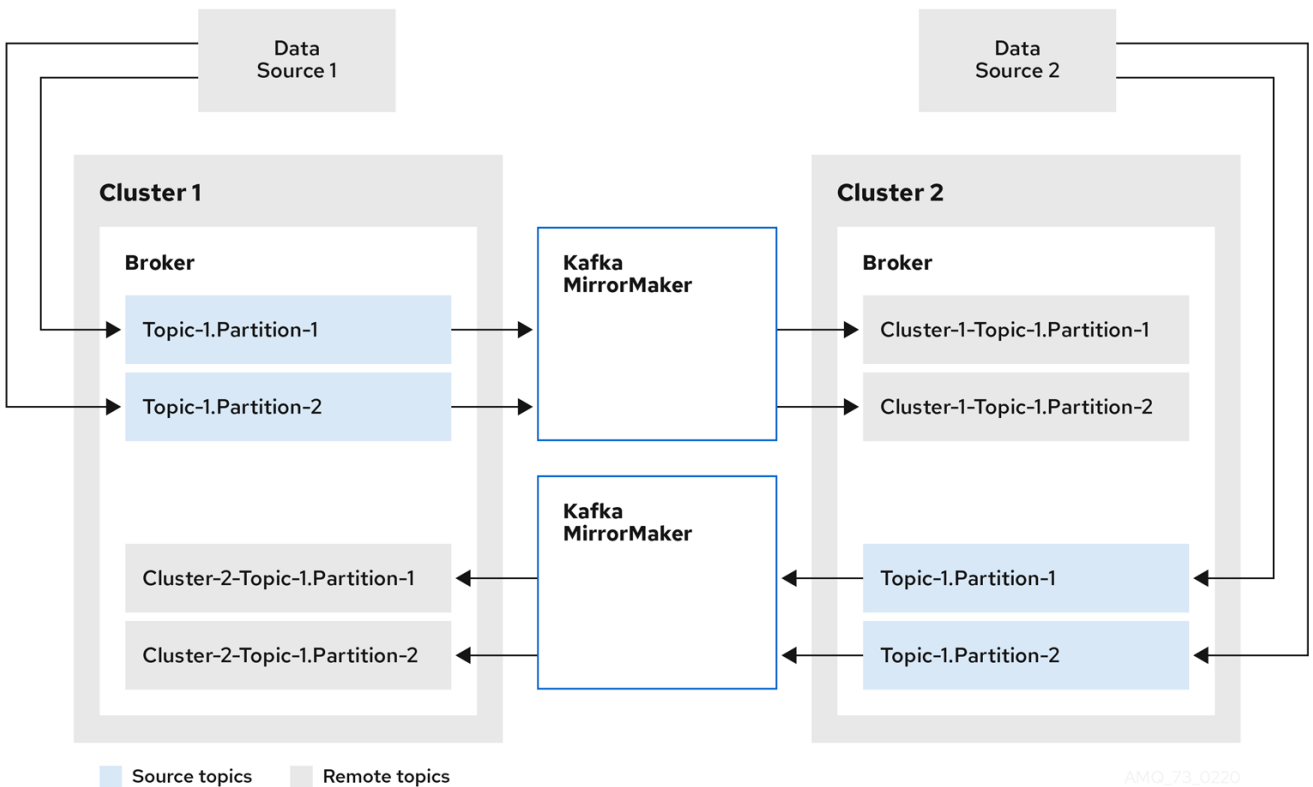
Bidirectional replication across two clusters

The MirrorMaker 2.0 architecture supports bidirectional replication in an *active/active* cluster configuration, so both clusters are active and provide the same data simultaneously. A MirrorMaker 2.0 cluster is required at each target destination.

Remote topics are distinguished by automatic renaming that prepends the name of cluster to the name of the topic. This is useful if you want to make the same data available locally in different geographical locations.

However, if you want to backup or migrate data in an active/passive cluster configuration, you might want to keep the original names of the topics. If so, you can configure MirrorMaker 2.0 to turn off automatic renaming.

Figure 5.2. Bidirectional replication



AMQ_73_0220

Example YAML showing MirrorMaker 2.0 configuration

```
apiVersion: kafka.strimzi.io/v1alpha1
```

```

kind: KafkaMirrorMaker2
metadata:
  name: my-mirror-maker2
spec:
  version: 2.6.0
  connectCluster: "my-cluster-target"
  clusters:
    - alias: "my-cluster-source"
      bootstrapServers: my-cluster-source-kafka-bootstrap:9092
    - alias: "my-cluster-target"
      bootstrapServers: my-cluster-target-kafka-bootstrap:9092
  mirrors:
    - sourceCluster: "my-cluster-source"
      targetCluster: "my-cluster-target"
      sourceConnector: {}
  topicsPattern: ".*"
  groupsPattern: "group1|group2|group3"

```

MirrorMaker

The earlier version of MirrorMaker uses producers and consumers to replicate data across clusters.

MirrorMaker uses:

- Consumer configuration to consume data from the source cluster
- Producer configuration to output data to the target cluster

Consumer and producer configuration includes any authentication and encryption settings.

A *whitelist* defines the topics to mirror from a source to a target cluster.

Key Consumer configuration

Consumer group identifier

The consumer group ID for a MirrorMaker consumer so that messages consumed are assigned to a consumer group.

Number of consumer streams

A value to determine the number of consumers in a consumer group that consume a message in parallel.

Offset commit interval

An offset commit interval to set the time between consuming and committing a message.

Key Producer configuration

Cancel option for send failure

You can define whether a message send failure is ignored or MirrorMaker is terminated and recreated.

Example YAML showing MirrorMaker configuration

```

apiVersion: kafka.strimzi.io/v1beta1
kind: KafkaMirrorMaker
metadata:
  name: my-mirror-maker
spec:

```

```
# ...
consumer:
  bootstrapServers: my-source-cluster-kafka-bootstrap:9092
  groupId: "my-group"
  numStreams: 2
  offsetCommitInterval: 120000
# ...
producer:
  # ...
  abortOnSendFailure: false
  # ...
whitelist: "my-topic|other-topic"
# ...
```

5.5. KAFKA CONNECT CONFIGURATION

A basic Kafka Connect configuration requires a bootstrap address to connect to a Kafka cluster, and encryption and authentication details.

Kafka Connect instances are configured by default with the same:

- Group ID for the Kafka Connect cluster
- Kafka topic to store the connector offsets
- Kafka topic to store connector and task status configurations
- Kafka topic to store connector and task status updates

If multiple different Kafka Connect instances are used, these settings must reflect each instance.

Example YAML showing Kafka Connect configuration

```
apiVersion: kafka.strimzi.io/v1beta1
kind: KafkaConnect
metadata:
  name: my-connect
spec:
  # ...
  config:
    group.id: my-connect-cluster
    offset.storage.topic: my-connect-cluster-offsets
    config.storage.topic: my-connect-cluster-configs
    status.storage.topic: my-connect-cluster-status
  # ...
```

Connectors

Connectors are configured separately from Kafka Connect. The configuration describes the source input data and target output data to feed into and out of Kafka Connect. The external source data must reference specific topics that will store the messages.

Kafka provides two built-in connectors:

- **FileStreamSourceConnector** streams data from an external system to Kafka, reading lines from an input source and sending each line to a Kafka topic.

- **FileStreamSinkConnector** streams data from Kafka to an external system, reading messages from a Kafka topic and creating a line for each in an output file.

You can add other connectors using connector plugins, which are a set of JAR files that define the implementation required to connect to certain types of external system.

You create a custom Kafka Connect image that uses new Kafka Connect plugins.

To create the image, you can use:

- A Kafka container image on [Red Hat Ecosystem Catalog](#) as a base image
- OpenShift [builds](#) and the [Source-to-Image \(S2I\)](#) framework to create new container images

Managing connectors

You can use the **KafkaConnector** resource or the [Kafka Connect REST API](#) to create and manage connector instances in a Kafka Connect cluster. The **KafkaConnector** resource offers an OpenShift-native approach, and is managed by the Cluster Operator.

The **spec** for the **KafkaConnector** resource specifies the connector class and configuration settings, as well as the maximum number of connector *tasks* to handle the data.

Example YAML showing KafkaConnector configuration

```
apiVersion: kafka.strimzi.io/v1beta1
kind: KafkaConnector
metadata:
  name: my-source-connector
  labels:
    strimzi.io/cluster: my-connect-cluster
spec:
  class: org.apache.kafka.connect.file.FileStreamSourceConnector
  tasksMax: 2
  config:
    file: "/opt/kafka/LICENSE"
    topic: my-topic
  # ...
```

You enable **KafkaConnectors** by adding an annotation to the **KafkaConnect** resource.

Example YAML showing annotation to enable KafkaConnector

```
apiVersion: kafka.strimzi.io/v1beta1
kind: KafkaConnect
metadata:
  name: my-connect
  annotations:
    strimzi.io/use-connector-resources: "true"
  # ...
```

5.6. KAFKA BRIDGE CONFIGURATION

A Kafka Bridge configuration requires a bootstrap server specification for the Kafka cluster it connects to, as well as any encryption and authentication options required.

Kafka Bridge consumer and producer configuration is standard, as described in the [Apache Kafka configuration documentation for consumers](#) and [Apache Kafka configuration documentation for producers](#).

HTTP-related configuration options set the port connection which the server listens on.

CORS

The Kafka Bridge supports the use of Cross-Origin Resource Sharing (CORS). CORS is a HTTP mechanism that allows browser access to selected resources from more than one origin, for example, resources on different domains. If you choose to use CORS, you can define a list of allowed resource origins and HTTP methods for interaction with the Kafka cluster through the Kafka Bridge. The lists are defined in the **http** specification of the Kafka Bridge configuration.

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

- A simple request is a HTTP request that must have an allowed origin defined in its header.
- A preflighted request sends an initial OPTIONS HTTP request before the actual request to check that the origin and the method are allowed.

Example YAML showing Kafka Bridge configuration

```
apiVersion: kafka.strimzi.io/v1alpha1
kind: KafkaBridge
metadata:
  name: my-bridge
spec:
  # ...
  bootstrapServers: my-cluster-kafka:9092
  http:
    port: 8080
    cors:
      allowedOrigins: "https://strimzi.io"
      allowedMethods: "GET,POST,PUT,DELETE,OPTIONS,PATCH"
  consumer:
    config:
      auto.offset.reset: earliest
  producer:
    config:
      delivery.timeout.ms: 300000
  # ...
```

Additional resources

- [Fetch](#) CORS specification

CHAPTER 6. SECURING KAFKA

A secure deployment of AMQ Streams can encompass:

- Encryption for data exchange
- Authentication to prove identity
- Authorization to allow or decline actions executed by users

6.1. ENCRYPTION

AMQ Streams supports Transport Layer Security (TLS), a protocol for encrypted communication.

Communication is always encrypted for communication between:

- Kafka brokers
- ZooKeeper nodes
- Operators and Kafka brokers
- Operators and ZooKeeper nodes
- Kafka Exporter

You can also configure TLS between Kafka brokers and clients by applying TLS encryption to the listeners of the Kafka broker. TLS is specified for external clients when configuring an external listener.

AMQ Streams components and Kafka clients use digital certificates for encryption. The Cluster Operator sets up certificates to enable encryption within the Kafka cluster. You can provide your own server certificates, referred to as *Kafka listener certificates*, for communication between Kafka clients and Kafka brokers, and inter-cluster communication.

AMQ Streams uses *Secrets* to store the certificates and private keys required for TLS in PEM and PKCS #12 format.

A TLS Certificate Authority (CA) issues certificates to authenticate the identity of a component. AMQ Streams verifies the certificates for the components against the CA certificate.

- AMQ Streams components are verified against the *cluster CA* Certificate Authority (CA)
- Kafka clients are verified against the *clients CA* Certificate Authority (CA)

6.2. AUTHENTICATION

Kafka listeners use authentication to ensure a secure client connection to the Kafka cluster.

Supported authentication mechanisms:

- Mutual TLS client authentication (on listeners with TLS enabled encryption)
- SASL SCRAM-SHA-512
- OAuth 2.0 token based authentication

The User Operator manages user credentials for TLS and SCRAM authentication, but not OAuth 2.0. For example, through the User Operator you can create a user representing a client that requires access to the Kafka cluster, and specify TLS as the authentication type.

Using OAuth 2.0 token-based authentication, application clients can access Kafka brokers without exposing account credentials. An authorization server handles the granting of access and inquiries about access.

6.3. AUTHORIZATION

Kafka clusters use authorization to control the operations that are permitted on Kafka brokers by specific clients or users. If applied to a Kafka cluster, authorization is enabled for all listeners used for client connection.

If a user is added to a list of *super users* in a Kafka broker configuration, the user is allowed unlimited access to the cluster regardless of any authorization constraints implemented through authorization mechanisms.

Supported authorization mechanisms:

- Simple authorization
- OAuth 2.0 authorization (if you are using OAuth 2.0 token-based authentication)
- Open Policy Agent (OPA) authorization

Simple authorization uses **AclAuthorizer**, the default Kafka authorization plugin. **AclAuthorizer** uses Access Control Lists (ACLs) to define which users have access to which resources.

OAuth 2.0 and OPA provide policy-based control from an authorization server. Security policies and permissions used to grant access to resources on Kafka brokers are defined in the authorization server.

URLs are used to connect to the authorization server and verify that an operation requested by a client or user is allowed or denied. Users and clients are matched against the policies created in the authorization server that permit access to perform specific actions on Kafka brokers.

CHAPTER 7. MONITORING

Monitoring data allows you to monitor the performance and health of AMQ Streams. You can configure your deployment to capture metrics data for analysis and notifications.

Metrics data is useful when investigating issues with connectivity and data delivery. For example, metrics data can identify under-replicated partitions or the rate at which messages are consumed. Alerting rules can provide time-critical notifications on such metrics through a specified communications channel. Monitoring visualizations present real-time metrics data to help determine when and how to update the configuration of your deployment. Example metrics configuration files are provided with AMQ Streams.

Distributed tracing complements the gathering of metrics data by providing a facility for end-to-end tracking of messages through AMQ Streams.

Cruise Control provides support for rebalancing of Kafka clusters, based on workload data.

Metrics and monitoring tools

AMQ Streams can employ the following tools for metrics and monitoring:

- **Prometheus** pulls metrics from Kafka, ZooKeeper and Kafka Connect clusters. The Prometheus **Alertmanager** plugin handles alerts and routes them to a notification service.
- **Kafka Exporter** adds additional Prometheus metrics
- **Grafana** provides dashboard visualizations of Prometheus metrics
- **Jaeger** provides distributed tracing support to track transactions between applications
- **Cruise Control** balances data across a Kafka cluster

Additional resources

- [Prometheus](#)
- [Kafka Exporter](#)
- [Grafana Labs](#)
- [Jaeger](#)
- [Cruise Control Wiki](#)

7.1. PROMETHEUS

Prometheus can extract metrics data from Kafka components and the AMQ Streams Operators.

To use Prometheus to obtain metrics data and provide alerts, Prometheus and the Prometheus Alertmanager plugin must be deployed. Kafka resources must also be deployed or redeployed with metrics configuration to expose the metrics data.

Prometheus scrapes the exposed metrics data for monitoring. Alertmanager issues alerts when conditions indicate potential problems, based on pre-defined alerting rules.

Sample metrics and alerting rules configuration files are provided with AMQ Streams. The sample alerting mechanism provided with AMQ Streams is configured to send notifications to a Slack channel.

7.2. GRAFANA

Grafana uses the metrics data exposed by Prometheus to present dashboard visualizations for monitoring.

A deployment of Grafana is required, with Prometheus added as a data source. Example dashboards, supplied with AMQ Streams as JSON files, are imported through the Grafana interface to present monitoring data.

7.3. KAFKA EXPORTER

Kafka Exporter is an open source project to enhance monitoring of Apache Kafka brokers and clients. Kafka Exporter is deployed with a Kafka cluster to extract additional Prometheus metrics data from Kafka brokers related to offsets, consumer groups, consumer lag, and topics. You can use the Grafana dashboard provided to visualize the data collected by Prometheus from Kafka Exporter.

A sample configuration file, alerting rules and Grafana dashboard for Kafka Exporter are provided with AMQ Streams.

7.4. DISTRIBUTED TRACING

Within a Kafka deployment, distributed tracing using Jaeger is supported for:

- MirrorMaker to trace messages from a source cluster to a target cluster
- Kafka Connect to trace messages consumed and produced by Kafka Connect
- Kafka Bridge to trace messages consumed and produced by Kafka Bridge, and HTTP requests from client applications

Template configuration properties are set for the Kafka resources, which describe tracing environment variables.

Tracing for Kafka clients

Client applications, such as Kafka producers and consumers, can also be set up so that transactions are monitored. Clients are configured with a tracing profile, and a tracer is initialized for the client application to use.

7.5. CRUISE CONTROL

Cruise Control is an open source project for simplifying the monitoring and balancing of data across a Kafka cluster. Cruise Control is deployed alongside a Kafka cluster to monitor its traffic, propose more balanced partition assignments, and trigger partition reassignments based on those proposals.

Cruise Control collects resource utilization information to model and analyze the workload of the Kafka cluster. Based on *optimization goals* that have been defined, Cruise Control generates *optimization proposals* outlining how the cluster can be effectively rebalanced. When an *optimization proposal* is approved, Cruise Control applies the rebalancing outlined in the proposal.

Prometheus can extract Cruise Control metrics data, including data related to optimization proposals and rebalancing operations. A sample configuration file and Grafana dashboard for Cruise Control are provided with AMQ Streams.

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 **Red Hat AMQ Streams** entries in the **JBOSS 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 2020-12-02 15:51:49 UTC