Red Hat AMQ 2021.q2

Deploying and Upgrading AMQ Streams on OpenShift

For use with AMQ Streams 1.7 on OpenShift Container Platform

Last Updated: 2021-06-04

Red Hat AMQ 2021.q2 Deploying and Upgrading AMQ Streams on OpenShift

For use with AMQ Streams 1.7 on OpenShift Container Platform

Legal Notice

Abstract

This guide provides instructions for deploying and upgrading AMQ Streams

Table of Contents

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. DEPLOYMENT OVERVIEW

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

This guide provides instructions on all the options available for deploying and upgrading AMQ Streams, describing what is deployed, and the order of deployment required to run Apache Kafka in an OpenShift cluster.

As well as describing the deployment steps, the guide also provides pre- and post-deployment instructions to prepare for and verify a deployment. Additional deployment options described include the steps to introduce metrics. Upgrade instructions are provided for AMQ Streams and Kafka upgrades.

AMQ Streams is designed to work on all types of OpenShift cluster regardless of distribution, from public and private clouds to local deployments intended for development.

1.1. 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

1.2. 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

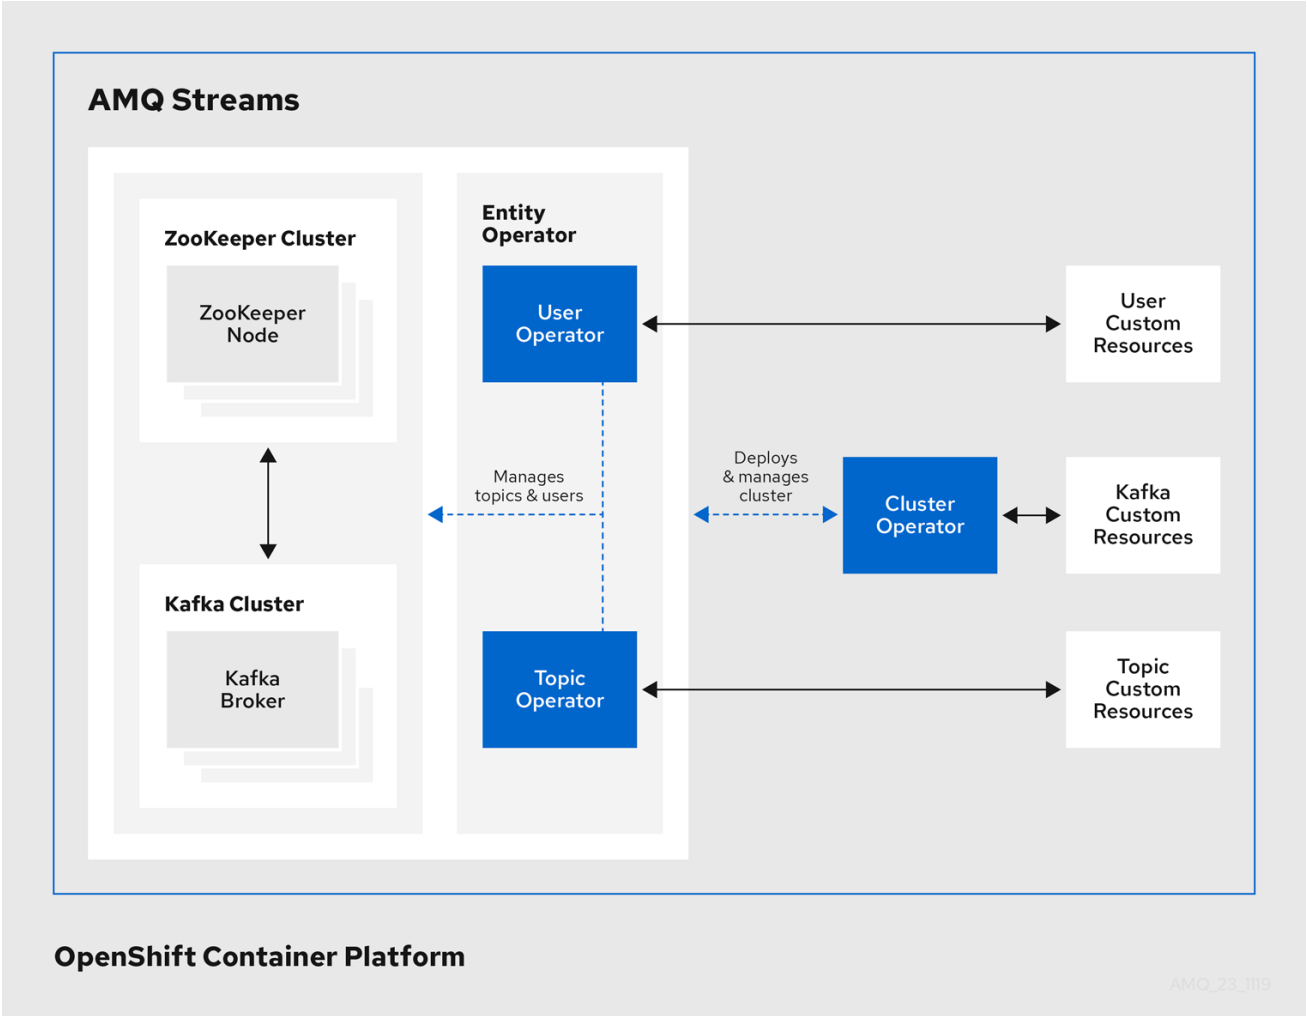

1.2.1. Cluster Operator

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

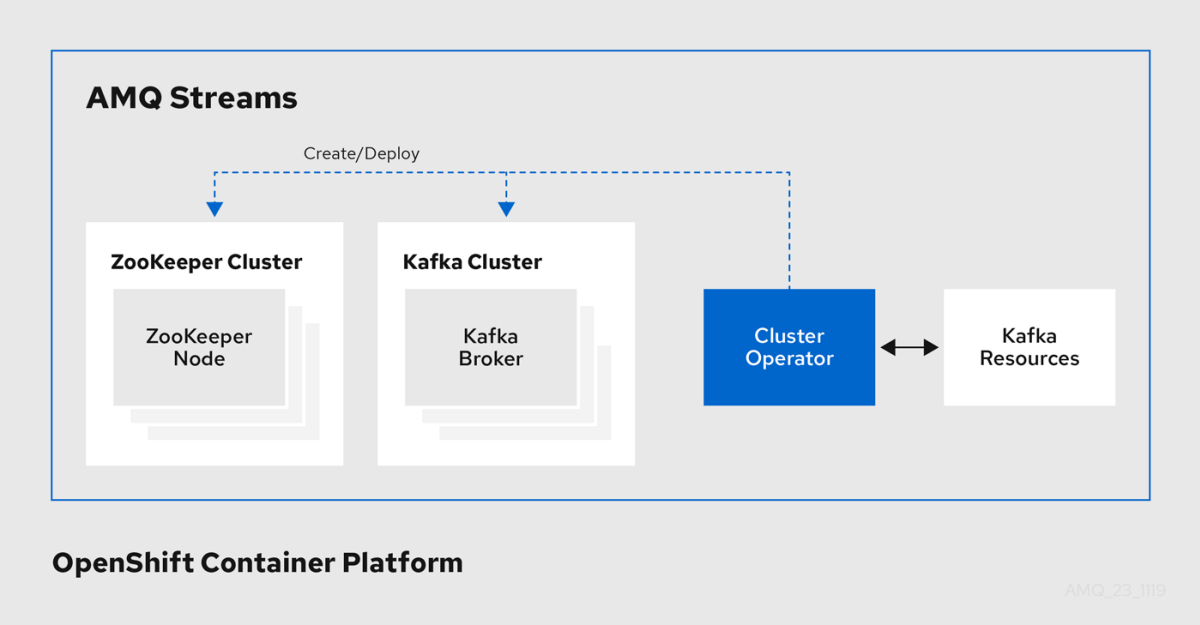

1.2.2. Topic Operator

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

Example architecture for the Topic Operator

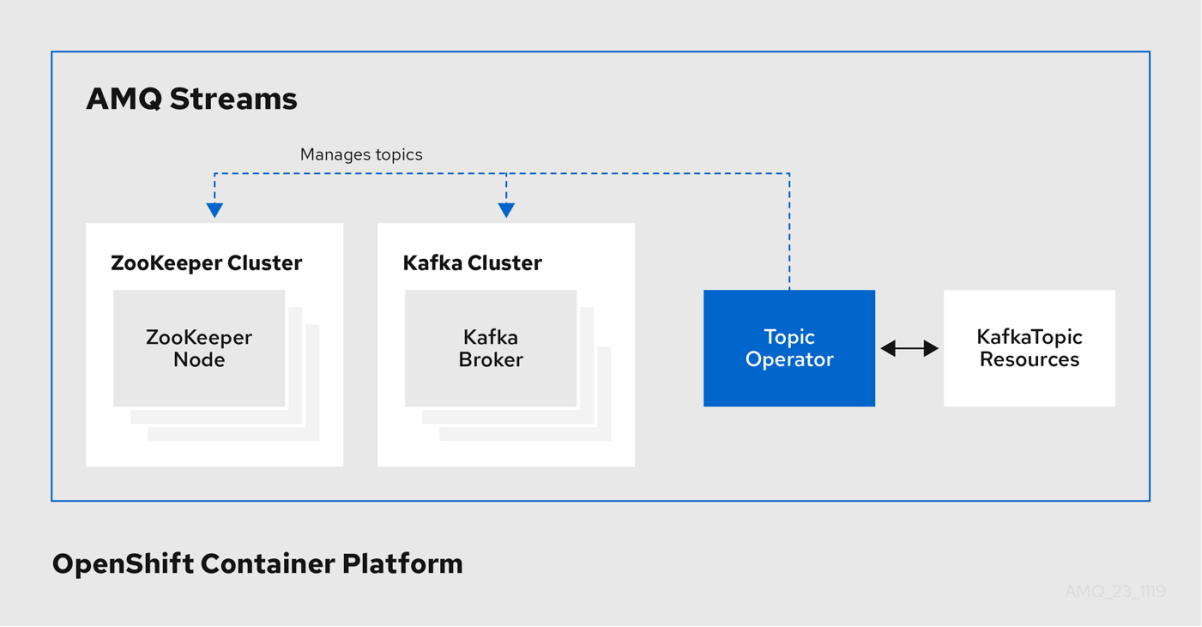

The role of the Topic Operator is to keep a set of **KafkaTopic** OpenShift resources describing Kafka topics in-sync with corresponding Kafka topics.

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.

The Topic Operator maintains information about each topic in a *topic store*, which is continually synchronized with updates from Kafka topics or OpenShift **KafkaTopic** custom resources. Updates from operations applied to a local in-memory topic store are persisted to a backup topic store on disk. If a topic is reconfigured or reassigned to other brokers, the **KafkaTopic** will always be up to date.

1.2.3. 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.

1.3. AMQ STREAMS CUSTOM RESOURCES

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.

Additional resources

- Extend the Kubernetes API with CustomResourceDefinitions

1.3.1. AMQ Streams custom resource example

CRDs require a one-time installation in a cluster to define the schemas used to instantiate and manage AMQ Streams-specific 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.

Depending on the cluster setup, installation typically requires cluster admin privileges.

NOTE

Access to manage custom resources is limited to AMQ Streams administrators. For more information, see Designating AMQ Streams administrators in the *Deploying and Upgrading AMQ Streams on OpenShift* guide.

A CRD defines a new **kind** of resource, such as **kind:Kafka**, within an OpenShift cluster.

The Kubernetes API server allows custom resources to be created based on the **kind** and understands from the CRD how to validate and store the custom resource when it is added to the OpenShift cluster.

WARNING

When CRDs are deleted, custom resources of that type are also deleted. Additionally, the resources created by the custom resource, such as pods and statefulsets are also deleted.

Each AMQ Streams-specific custom resource conforms to the schema defined by the CRD for the resource's **kind**. The custom resources for AMQ Streams components have common configuration properties, which are defined under **spec**.

To understand the relationship between a CRD and a custom resource, let's look at a sample of the CRD for a Kafka topic.

Kafka topic CRD

```
apiVersion: kafka.strimzi.io/v1beta2
kind: CustomResourceDefinition
metadata: 1
  name: kafkatopics.kafka.strimzi.io
  labels:
    app: strimzi
spec: 2
  group: kafka.strimzi.io
  versions:
    v1beta2
  scope: Namespaced
  names:
    # ...
    singular: kafkatopic
    plural: kafkatopics
    shortNames:
    - kt 3
  additionalPrinterColumns: 4
      # ...
  subresources:
    status: {} 5
  validation: 6
    openAPIV3Schema:
      properties:
        spec:
          type: object
          properties:
            partitions:
              type: integer
              minimum: 1
            replicas:
              type: integer
              minimum: 1
              maximum: 32767
      # ...
```

1. The metadata for the topic CRD, its name and a label to identify the CRD.
2. The specification for this CRD, including the group (domain) name, the plural name and the supported schema version, which are used in the URL to access the API of the topic. The other names are used to identify instance resources in the CLI. For example, **oc get kafkatopic my-topic** or **oc get kafkatopics**.
3. The shortname can be used in CLI commands. For example, **oc get kt** can be used as an abbreviation instead of **oc get kafkatopic**.
4. The information presented when using a **get** command on the custom resource.
5. The current status of the CRD as described in the schema reference for the resource.
6. openAPIV3Schema validation provides validation for the creation of topic custom resources. For example, a topic requires at least one partition and one replica.

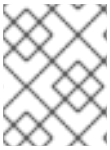

NOTE

You can identify the CRD YAML files supplied with the AMQ Streams installation files, because the file names contain an index number followed by 'Crd'.

Here is a corresponding example of a **KafkaTopic** custom resource.

Kafka topic custom resource

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaTopic (1)
metadata:
  name: my-topic
  labels:
    strimzi.io/cluster: my-cluster (2)
spec: (3)
  partitions: 1
  replicas: 1
  config:
    retention.ms: 7200000
    segment.bytes: 1073741824
status:
  conditions: (4)
    lastTransitionTime: "2019-08-20T11:37:00.706Z"
    status: "True"
    type: Ready
  observedGeneration: 1
  / ...
```

1. The **kind** and **apiVersion** identify the CRD of which the custom resource is an instance.
2. A label, applicable only to **KafkaTopic** and **KafkaUser** resources, that defines the name of the Kafka cluster (which is same as the name of the **Kafka** resource) to which a topic or user belongs.
3. The spec shows the number of partitions and replicas for the topic as well as the configuration parameters for the topic itself. In this example, the retention period for a message to remain in the topic and the segment file size for the log are specified.

4 Status conditions for the **KafkaTopic** resource. The **type** condition changed to **Ready** at the **lastTransitionTime**.

Custom resources can be applied to a cluster through the platform CLI. When the custom resource is created, it uses the same validation as the built-in resources of the Kubernetes API.

After a **KafkaTopic** custom resource is created, the Topic Operator is notified and corresponding Kafka topics are created in AMQ Streams.

1.4. AMQ STREAMS INSTALLATION METHODS

There are two ways to install AMQ Streams on OpenShift.

Installation method	Description	Supported platform
Installation artifacts (YAML files)	Download the **amq-streams-x.y.z-ocp-install-examples.zip** file from the AMQ Streams download site. Next, deploy the YAML installation artifacts to your OpenShift cluster using **oc**. You start by deploying the Cluster Operator from **install/cluster-operator** to a single namespace, multiple namespaces, or all namespaces.	OpenShift 4.6 and later
OperatorHub	Use the **Red Hat Integration - AMQ Streams Operator** in the OperatorHub to deploy AMQ Streams to a single namespace or all namespaces.	OpenShift 4.6 and later

For the greatest flexibility, choose the installation artifacts method. Choose the OperatorHub method if you want to install AMQ Streams to OpenShift 4.6 and later in a standard configuration using the web console. The OperatorHub also allows you to take advantage of automatic updates.

Both methods install the Cluster Operator to your OpenShift cluster. Use the *same* method to deploy the other components, starting with the Kafka cluster. If you are using the installation artifacts method, example YAML files are provided. If you are using the OperatorHub, the AMQ Streams Operator makes Kafka components available for installation from the OpenShift web console.

AMQ Streams installation artifacts

The AMQ Streams installation artifacts contain various YAML files that can be deployed to OpenShift, using **oc**, to create custom resources, including:

- Deployments
- Custom resource definitions (CRDs)
- Roles and role bindings
- Service accounts

YAML installation files are provided for the Cluster Operator, Topic Operator, User Operator, and the Strimzi Admin role.

OperatorHub

Starting with OpenShift 4, the *Operator Lifecycle Manager (OLM)* helps cluster administrators to install, update, and manage the lifecycle of all Operators and their associated services running across their clusters. The OLM is part of the *Operator Framework*, an open source toolkit designed to manage Kubernetes-native applications (Operators) in an effective, automated, and scalable way.

The *OperatorHub* is part of the OpenShift web console. Cluster administrators can use it to discover, install, and upgrade Operators. Operators can be pulled from the OperatorHub, installed on the OpenShift cluster to a single namespace or all namespaces, and managed by the OLM. Engineering teams can then independently manage the software in development, test, and production environments using the OLM.

Red Hat Integration - AMQ Streams Operator

The *Red Hat Integration - AMQ Streams Operator* is available to install from the OperatorHub. Once installed, the AMQ Streams Operator deploys the Cluster Operator to your OpenShift cluster, along with the necessary CRDs and role-based access control (RBAC) resources. You still need to install the Kafka components from the OpenShift web console.

Additional resources

Installing AMQ Streams using the installation artifacts:

- Section 5.1.1.2, “Deploying the Cluster Operator to watch a single namespace”
- Section 5.1.1.3, “Deploying the Cluster Operator to watch multiple namespaces”
- Section 5.1.1.4, “Deploying the Cluster Operator to watch all namespaces”

Installing AMQ Streams from the OperatorHub:

- Section 4.2, “Deploying the AMQ Streams Operator from the OperatorHub”
- *Operators* guide in the OpenShift documentation.

CHAPTER 2. WHAT IS DEPLOYED WITH AMQ STREAMS

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.

2.1. ORDER OF DEPLOYMENT

The required order of deployment to an OpenShift cluster is as follows:

1. Deploy the Cluster operator to manage your Kafka cluster
2. Deploy the Kafka cluster with the ZooKeeper cluster, and include the Topic Operator and User Operator in the deployment
3. Optionally deploy:
    - The Topic Operator and User Operator standalone if you did not deploy them with the Kafka cluster
    - Kafka Connect
    - Kafka MirrorMaker
    - Kafka Bridge
    - Components for the monitoring of metrics

2.2. ADDITIONAL DEPLOYMENT CONFIGURATION OPTIONS

The deployment procedures in this guide describe a deployment using the example installation YAML files provided with AMQ Streams. The procedures highlight any important configuration considerations, but they do not describe all the configuration options available.

You can use custom resources to refine your deployment.

You may wish to review the configuration options available for Kafka components before you deploy AMQ Streams. For more information on the configuration through custom resources, see Deployment configuration in the *Using AMQ Streams on OpenShift* guide.

2.2.1. Securing Kafka

On deployment, the Cluster Operator automatically sets up TLS certificates for data encryption and authentication within your cluster.

AMQ Streams provides additional configuration options for *encryption*, *authentication* and *authorization*, which are described in the *Using AMQ Streams on OpenShift* guide:

- Secure data exchange between the Kafka cluster and clients by Managing secure access to Kafka.
- Configure your deployment to use an authorization server to provide OAuth 2.0 authentication and OAuth 2.0 authorization.
- Secure Kafka using your own certificates .

2.2.2. Monitoring your deployment

AMQ Streams supports additional deployment options to monitor your deployment.

- Extract metrics and monitor Kafka components by deploying Prometheus and Grafana with your Kafka cluster.
- Extract additional metrics, particularly related to monitoring consumer lag, by deploying Kafka Exporter with your Kafka cluster.
- Track messages end-to-end by setting up distributed tracing, as described in the *Using AMQ Streams on OpenShift* guide.

CHAPTER 3. PREPARING FOR YOUR AMQ STREAMS DEPLOYMENT

This section shows how you prepare for a AMQ Streams deployment, describing:

- The prerequisites you need before you can deploy AMQ Streams
- How to download the AMQ Streams release artifacts to use in your deployment
- How to authenticate with the Red Hat registry for Kafka Connect Source-to-Image (S2I) builds (if required)
- How to push the AMQ Streams container images into your own registry (if required)
- How to set up *admin* roles for configuration of custom resources used in deployment

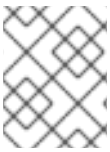

NOTE

To run the commands in this guide, your cluster user must have the rights to manage role-based access control (RBAC) and CRDs.

3.1. DEPLOYMENT PREREQUISITES

To deploy AMQ Streams, make sure:

- An OpenShift 4.6 and later cluster is available
  AMQ Streams is based on AMQ Streams Strimzi 0.22.x.
- The **oc** command-line tool is installed and configured to connect to the running cluster.

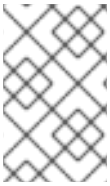

NOTE

AMQ Streams supports some features that are specific to OpenShift, where such integration benefits OpenShift users and there is no equivalent implementation using standard OpenShift.

3.2. DOWNLOADING AMQ STREAMS RELEASE ARTIFACTS

To install AMQ Streams, download and extract the release artifacts from the **amq-streams-*<version>*-ocp-install-examples.zip** file from the AMQ Streams download site .

AMQ Streams release artifacts include sample YAML files to help you deploy the components of AMQ Streams to OpenShift, perform common operations, and configure your Kafka cluster.

Use **oc** to deploy the Cluster Operator from the **install/cluster-operator** folder of the downloaded ZIP file. For more information about deploying and configuring the Cluster Operator, see Section 5.1.1, “Deploying the Cluster Operator”.

In addition, if you want to use standalone installations of the Topic and User Operators with a Kafka cluster that is not managed by the AMQ Streams Cluster Operator, you can deploy them from the **install/topic-operator** and **install/user-operator** folders.

NOTE

Additionally, AMQ Streams container images are available through the Red Hat Ecosystem Catalog. However, we recommend that you use the YAML files provided to deploy AMQ Streams.

3.3. AUTHENTICATING WITH THE CONTAINER REGISTRY FOR KAFKA CONNECT S2I

You need to configure authentication with the Red Hat container registry (**registry.redhat.io**) before creating a container image using OpenShift builds and Source-to-Image (S2I) .

The container registry is used to store AMQ Streams container images on the Red Hat Ecosystem Catalog. The Catalog contains a Kafka Connect builder image with S2I support. The OpenShift build pulls this builder image, together with your source code and binaries, and uses it to build the new container image.

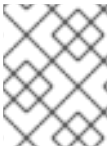

NOTE

Authentication with the Red Hat container registry is only required if using Kafka Connect S2I. It is not required for the other AMQ Streams components.

Prerequisites

- Cluster administrator access to an OpenShift Container Platform cluster.
- Login details for your Red Hat Customer Portal account. See Appendix A, *Using your subscription*.

Procedure

1. If needed, log in to your OpenShift cluster as an administrator:

```
oc login --user system:admin --token=my-token --server=https://my-cluster.example.com:6443
```

2. Open the project that will contain the Kafka Connect S2I cluster:

```
oc project CLUSTER-NAME
```

NOTE

You might have already deployed the Kafka Connect S2I cluster .

3. Create a **docker-registry** secret using your Red Hat Customer Portal account, replacing **PULL-SECRET-NAME** with the secret name to create:

```
oc create secret docker-registry PULL-SECRET-NAME \
 --docker-server=registry.redhat.io \
 --docker-username=CUSTOMER-PORTAL-USERNAME \
 --docker-password=CUSTOMER-PORTAL-PASSWORD \
 --docker-email=EMAIL-ADDRESS
```

You should see the following output:

```
secret/PULL-SECRET-NAME created
```

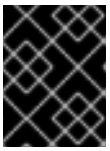

IMPORTANT

You must create this **docker-registry** secret in every OpenShift project that will authenticate to **registry.redhat.io**.

4. Link the secret to your service account to use the secret for pulling images. The service account name must match the name that the OpenShift pod uses.

```
oc secrets link SERVICE-ACCOUNT-NAME PULL-SECRET-NAME --for=pull
```

For example, using the **default** service account and a secret named **my-secret**:

```
oc secrets link default my-secret --for=pull
```

5. Link the secret to the **builder** service account to use the secret for pushing and pulling build images:

```
oc secrets link builder PULL-SECRET-NAME
```

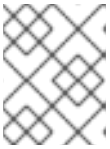

NOTE

If you do not want to use your Red Hat username and password to create the pull secret, you can create an authentication token using a registry service account.

Additional resources

- Section 5.2.3.3, "Creating a container image using OpenShift builds and Source-to-Image"
- Red Hat Container Registry authentication (Red Hat Knowledgebase)
- Registry Service Accounts on the Red Hat Customer Portal

3.4. PUSHING CONTAINER IMAGES TO YOUR OWN REGISTRY

Container images for AMQ Streams are available in the Red Hat Ecosystem Catalog. The installation YAML files provided by AMQ Streams will pull the images directly from the Red Hat Ecosystem Catalog.

If you do not have access to the Red Hat Ecosystem Catalog or want to use your own container repository:

1. Pull **all** container images listed here
2. Push them into your own registry
3. Update the image names in the installation YAML files

NOTE

Each Kafka version supported for the release has a separate image.

Container image	Namespace/Repository	Description
Kafka	• registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0 • registry.redhat.io/amq7/amq-streams-kafka-26-rhel7:1.7.0	AMQ Streams image for running Kafka, including: • Kafka Broker • Kafka Connect / S2I • Kafka Mirror Maker • ZooKeeper • TLS Sidecars
Operator	• registry.redhat.io/amq7/amq-streams-rhel7-operator:1.7.0	AMQ Streams image for running the operators: • Cluster Operator • Topic Operator • User Operator • Kafka Initializer
Kafka Bridge	• registry.redhat.io/amq7/amq-streams-bridge-rhel7:1.7.0	AMQ Streams image for running the AMQ Streams Kafka Bridge

3.5. DESIGNATING AMQ STREAMS ADMINISTRATORS

AMQ Streams provides custom resources for configuration of your deployment. By default, permission to view, create, edit, and delete these resources is limited to OpenShift cluster administrators. AMQ Streams provides two cluster roles that you can use to assign these rights to other users:

- **strimzi-view** allows users to view and list AMQ Streams resources.
- **strimzi-admin** allows users to also create, edit or delete AMQ Streams resources.

When you install these roles, they will automatically aggregate (add) these rights to the default OpenShift cluster roles. **strimzi-view** aggregates to the **view** role, and **strimzi-admin** aggregates to the **edit** and **admin** roles. Because of the aggregation, you might not need to assign these roles to users who already have similar rights.

The following procedure shows how to assign a **strimzi-admin** role that allows non-cluster administrators to manage AMQ Streams resources.

A system administrator can designate AMQ Streams administrators after the Cluster Operator is deployed.

Prerequisites

- The AMQ Streams Custom Resource Definitions (CRDs) and role-based access control (RBAC) resources to manage the CRDs have been deployed with the Cluster Operator.

Procedure

1. Create the **strimzi-view** and **strimzi-admin** cluster roles in OpenShift.

```
oc create -f install/strimzi-admin
```

2. If needed, assign the roles that provide access rights to users that require them.

```
oc create clusterrolebinding strimzi-admin --clusterrole=strimzi-admin --user=user1 --user=user2
```

CHAPTER 4. DEPLOYING AMQ STREAMS FROM THE OPERATORHUB

Use the Red Hat Integration - AMQ Streams Operator to deploy AMQ Streams from the OperatorHub.

The procedures in this section show how to:

- Deploy the AMQ Streams Operator from the OperatorHub
- Deploy Kafka components using the AMQ Streams Operator

4.1. USING THE RED HAT INTEGRATION OPERATOR TO INSTALL THE AMQ STREAMS OPERATOR

The Red Hat Integration Operator allows you to choose and install the Operators that manage your Red Hat Integration components. If you have more than one Red Hat Integration subscription, you can use the Red Hat Integration Operator to install and update the AMQ Streams Operator, as well as the Operators for all subscribed Red Hat Integration components.

As with the AMQ Streams Operator, you can use the Operator Lifecycle Manager (OLM) to install the Red Hat Integration Operator on an OpenShift Container Platform (OCP) cluster from the OperatorHub in the OCP console.

Additional resources

For more information on installing and using the Red Hat Integration Operator, see Installing the Red Hat Integration Operator on OpenShift.

4.2. DEPLOYING THE AMQ STREAMS OPERATOR FROM THE OPERATORHUB

You can deploy the Cluster Operator to your OpenShift cluster by installing the AMQ Streams Operator from the OperatorHub.

WARNING

Make sure you use the appropriate update channel. If you are on a supported version of the OpenShift, installing AMQ Streams from the default **stable** channel is *safe*. However, if you are using a version of the OpenShift that is unsupported, installing AMQ Streams from the stable channel is *unsafe*, especially when automatic updates are enabled, as the cluster will receive automatic updates with new components that are unsupported by the OpenShift release.

Prerequisites

- The **Red Hat Operators OperatorSource** is enabled in your OpenShift cluster. If you can see Red Hat Operators in the OperatorHub, the correct **OperatorSource** is enabled. For more information, see the *Operators* guide.

- Installation requires a user with sufficient privileges to install Operators from the OperatorHub.

Procedure

1. In the OpenShift web console, click **Operators > OperatorHub**.
2. Search or browse for the **AMQ Streams** Operator, in the **Streaming & Messaging** category.
3. Click the **Red Hat Integration - AMQ Streams Operator** tile and then, in the sidebar on the right, click **Install**.
4. On the **Create Operator Subscription** screen, choose from the following installation and update options:
   - **Update Channel**: Choose the update channel for the AMQ Streams Operator.
     - The (default) **stable** channel contains all the latest updates and releases, including major, minor, and micro releases, which are assumed to be well tested and stable.
     - An **amq-streams-*X*.x** channel contains the minor and micro release updates for a major release, where *X* is the major release version number.
     - An **amq-streams-*X.Y*.x** channel contains the micro release updates for a minor release, where *X* is the major release version number and *Y* is the minor release version number.
   - **Installation Mode**: Choose to install the AMQ Streams Operator to all namespaces in the cluster (the default option) or a specific namespace. It is good practice to use namespaces to separate functions. We recommend that you dedicate a specific namespace to the Kafka cluster and other AMQ Streams components.
   - **Approval Strategy**: By default, the AMQ Streams Operator is automatically upgraded to the latest AMQ Streams version by the Operator Lifecycle Manager (OLM). Optionally, select **Manual** if you want to manually approve future upgrades. For more information, see the *Operators* guide in the OpenShift documentation.
5. Click **Subscribe**; the AMQ Streams Operator is installed to your OpenShift cluster.
   The AMQ Streams Operator deploys the Cluster Operator, CRDs, and role-based access control (RBAC) resources to the selected namespace, or to all namespaces.
6. On the **Installed Operators** screen, check the progress of the installation. The AMQ Streams Operator is ready to use when its status changes to **InstallSucceeded**.

Next, you can use the AMQ Streams Operator to deploy the Kafka components, starting with a Kafka cluster.

Additional resources

- Section 4.3, "Deploying Kafka components using the AMQ Streams Operator"
- Section 1.4, "AMQ Streams installation methods"
- Section 5.1.2.1, "Deploying the Kafka cluster"

4.3. DEPLOYING KAFKA COMPONENTS USING THE AMQ STREAMS OPERATOR

When installed on an Openshift Container Platform, the AMQ Streams Operator makes Kafka components available for installation from the user interface.

Kafka components available for installation:

- Kafka
- Kafka Connect
- Kafka Connect Source to Image (S2I)
- Kafka MirrorMaker
- Kafka MirrorMaker 2
- Kafka Topic
- Kafka User
- Kafka Bridge
- Kafka Connector
- Kafka Rebalance

Prerequisites

- AMQ Streams Operator is installed on the OpenShift Container Platform (OCP) cluster

Procedure

1. Navigate to **Installed Operators** and click **Red Hat Integration - AMQ Streams Operator**to display the **Operator details** page.
2. From **Provided APIs**, click **Create Instance** for the Kafka component you wish to install. The default configuration for each component is encapsulated in a CRD **spec** property.
3. (Optional) Configure the installation specification from the *form* or *YAML* views before you perform the installation.
4. Click **Create** to start the installation of the selected component. Wait until the status changes to **Succeeded**.

Additional resources

- Section 4.2, “Deploying the AMQ Streams Operator from the OperatorHub”

CHAPTER 5. DEPLOYING AMQ STREAMS USING INSTALLATION ARTIFACTS

As an alternative to using the OperatorHub to deploy AMQ Streams using the AMQ Streams Operator, you can use the installation artifacts. Having prepared your environment for a deployment of AMQ Streams, this section shows:

- How to create the Kafka cluster
- Optional procedures to deploy other Kafka components according to your requirements:
  - Kafka Connect
  - Kafka MirrorMaker
  - Kafka Bridge

The procedures assume an OpenShift cluster is available and running.

AMQ Streams is based on AMQ Streams Strimzi 0.22.x. This section describes the procedures to deploy AMQ Streams on OpenShift 4.6 and later.

NOTE

To run the commands in this guide, your cluster user must have the rights to manage role-based access control (RBAC) and CRDs.

5.1. CREATE THE KAFKA CLUSTER

In order to create your Kafka cluster, you deploy the Cluster Operator to manage the Kafka cluster, then deploy the Kafka cluster.

When deploying the Kafka cluster using the **Kafka** resource, you can deploy the Topic Operator and User Operator at the same time. Alternatively, if you are using a non-AMQ Streams Kafka cluster, you can deploy the Topic Operator and User Operator as standalone components.

Deploying a Kafka cluster with the Topic Operator and User Operator

Perform these deployment steps if you want to use the Topic Operator and User Operator with a Kafka cluster managed by AMQ Streams.

1. Deploy the Cluster Operator
2. Use the Cluster Operator to deploy the:
   a. Kafka cluster
   b. Topic Operator
   c. User Operator

Deploying a standalone Topic Operator and User Operator

Perform these deployment steps if you want to use the Topic Operator and User Operator with a Kafka cluster that is **not managed** by AMQ Streams.

1. Deploy the standalone Topic Operator

2. Deploy the standalone User Operator

5.1.1. Deploying the Cluster Operator

The Cluster Operator is responsible for deploying and managing Apache Kafka clusters within an OpenShift cluster.

The procedures in this section show:

- How to deploy the Cluster Operator to *watch*:
  - A single namespace
  - Multiple namespaces
  - All namespaces
- Alternative deployment options:

5.1.1.1. Watch options for a Cluster Operator deployment

When the Cluster Operator is running, it starts to *watch* for updates of Kafka resources.

You can choose to deploy the Cluster Operator to watch Kafka resources from:

- A single namespace (the same namespace containing the Cluster Operator)
- Multiple namespaces
- All namespaces

NOTE

AMQ Streams provides example YAML files to make the deployment process easier.

The Cluster Operator watches for changes to the following resources:

- **Kafka** for the Kafka cluster.
- **KafkaConnect** for the Kafka Connect cluster.
- **KafkaConnectS2I** for the Kafka Connect cluster with Source2Image support.
- **KafkaConnector** for creating and managing connectors in a Kafka Connect cluster.
- **KafkaMirrorMaker** for the Kafka MirrorMaker instance.
- **KafkaBridge** for the Kafka Bridge instance

When one of these resources is created in the OpenShift cluster, the operator gets the cluster description from the resource and starts creating a new cluster for the resource by creating the necessary OpenShift resources, such as StatefulSets, Services and ConfigMaps.

Each time a Kafka resource is updated, the operator performs corresponding updates on the OpenShift resources that make up the cluster for the resource.

Resources are either patched or deleted, and then recreated in order to make the cluster for the resource reflect the desired state of the cluster. This operation might cause a rolling update that might lead to service disruption.

When a resource is deleted, the operator undeploys the cluster and deletes all related OpenShift resources.

5.1.1.2. Deploying the Cluster Operator to watch a single namespace

This procedure shows how to deploy the Cluster Operator to watch AMQ Streams resources in a single namespace in your OpenShift cluster.

Prerequisites

- This procedure requires use of an OpenShift user account which is able to create **CustomResourceDefinitions**, **ClusterRoles** and **ClusterRoleBindings**. Use of Role Base Access Control (RBAC) in the OpenShift cluster usually means that permission to create, edit, and delete these resources is limited to OpenShift cluster administrators, such as **system:admin**.

Procedure

1. Edit the AMQ Streams installation files to use the namespace the Cluster Operator is going to be installed into.
   For example, in this procedure the Cluster Operator is installed into the namespace ***my-cluster-operator-namespace***.

   On Linux, use:

   ```
   sed -i 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

   On MacOS, use:

   ```
   sed -i '' 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

2. Deploy the Cluster Operator:

   ```
   oc create -f install/cluster-operator -n my-cluster-operator-namespace
   ```

3. Verify that the Cluster Operator was successfully deployed:

   ```
   oc get deployments
   ```

5.1.1.3. Deploying the Cluster Operator to watch multiple namespaces

This procedure shows how to deploy the Cluster Operator to watch AMQ Streams resources across multiple namespaces in your OpenShift cluster.

Prerequisites

- This procedure requires use of an OpenShift user account which is able to create

CustomResourceDefinitions, **ClusterRoles** and **ClusterRoleBindings**. Use of Role Base Access Control (RBAC) in the OpenShift cluster usually means that permission to create, edit, and delete these resources is limited to OpenShift cluster administrators, such as **system:admin**.

Procedure

1. Edit the AMQ Streams installation files to use the namespace the Cluster Operator is going to be installed into.
   For example, in this procedure the Cluster Operator is installed into the namespace ***my-cluster-operator-namespace***.

   On Linux, use:

   ```
   sed -i 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

   On MacOS, use:

   ```
   sed -i '' 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

2. Edit the **install/cluster-operator/060-Deployment-strimzi-cluster-operator.yaml** file to add a list of all the namespaces the Cluster Operator will watch to the **STRIMZI_NAMESPACE** environment variable.
   For example, in this procedure the Cluster Operator will watch the namespaces **watched-namespace-1**, **watched-namespace-2**, **watched-namespace-3**.

   ```
   apiVersion: apps/v1
   kind: Deployment
   spec:
     # ...
     template:
       spec:
         serviceAccountName: strimzi-cluster-operator
         containers:
         - name: strimzi-cluster-operator
           image: registry.redhat.io/amq7/amq-streams-rhel7-operator:1.7.0
           imagePullPolicy: IfNotPresent
           env:
           - name: STRIMZI_NAMESPACE
             value: watched-namespace-1,watched-namespace-2,watched-namespace-3
   ```

3. For each namespace listed, install the **RoleBindings**.
   In this example, we replace ***watched-namespace*** in these commands with the namespaces listed in the previous step, repeating them for **watched-namespace-1**, **watched-namespace-2**, **watched-namespace-3**:

   ```
   oc create -f install/cluster-operator/020-RoleBinding-strimzi-cluster-operator.yaml -n watched-namespace
   oc create -f install/cluster-operator/031-RoleBinding-strimzi-cluster-operator-entity-operator-delegation.yaml -n watched-namespace
   oc create -f install/cluster-operator/032-RoleBinding-strimzi-cluster-operator-topic-operator-delegation.yaml -n watched-namespace
   ```

4. Deploy the Cluster Operator:

```
oc create -f install/cluster-operator -n my-cluster-operator-namespace
```

5. Verify that the Cluster Operator was successfully deployed:

```
oc get deployments
```

5.1.1.4. Deploying the Cluster Operator to watch all namespaces

This procedure shows how to deploy the Cluster Operator to watch AMQ Streams resources across all namespaces in your OpenShift cluster.

When running in this mode, the Cluster Operator automatically manages clusters in any new namespaces that are created.

Prerequisites

- This procedure requires use of an OpenShift user account which is able to create **CustomResourceDefinitions**, **ClusterRoles** and **ClusterRoleBindings**. Use of Role Base Access Control (RBAC) in the OpenShift cluster usually means that permission to create, edit, and delete these resources is limited to OpenShift cluster administrators, such as **system:admin**.

Procedure

1. Edit the AMQ Streams installation files to use the namespace the Cluster Operator is going to be installed into.
   For example, in this procedure the Cluster Operator is installed into the namespace ***my-cluster-operator-namespace***.

   On Linux, use:

```
sed -i 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
```

   On MacOS, use:

```
sed -i '' 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
```

2. Edit the **install/cluster-operator/060-Deployment-strimzi-cluster-operator.yaml** file to set the value of the **STRIMZI_NAMESPACE** environment variable to ***.

```
apiVersion: apps/v1
kind: Deployment
spec:
  # ...
  template:
    spec:
      # ...
      serviceAccountName: strimzi-cluster-operator
      containers:
```

```
      - name: strimzi-cluster-operator
        image: registry.redhat.io/amq7/amq-streams-rhel7-operator:1.7.0
        imagePullPolicy: IfNotPresent
        env:
        - name: STRIMZI_NAMESPACE
          value: "*"
        # ...
```

3. Create **ClusterRoleBindings** that grant cluster-wide access for all namespaces to the Cluster Operator.

```
oc create clusterrolebinding strimzi-cluster-operator-namespaced --clusterrole=strimzi-cluster-operator-namespaced --serviceaccount my-cluster-operator-namespace:strimzi-cluster-operator
oc create clusterrolebinding strimzi-cluster-operator-entity-operator-delegation --clusterrole=strimzi-entity-operator --serviceaccount my-cluster-operator-namespace:strimzi-cluster-operator
oc create clusterrolebinding strimzi-cluster-operator-topic-operator-delegation --clusterrole=strimzi-topic-operator --serviceaccount my-cluster-operator-namespace:strimzi-cluster-operator
```

   Replace ***my-cluster-operator-namespace*** with the namespace you want to install the Cluster Operator into.

4. Deploy the Cluster Operator to your OpenShift cluster.

```
oc create -f install/cluster-operator -n my-cluster-operator-namespace
```

5. Verify that the Cluster Operator was successfully deployed:

```
oc get deployments
```

5.1.2. Deploying Kafka

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

The procedures in this section show:

- How to use the Cluster Operator to deploy:
  - An *ephemeral* or *persistent* Kafka cluster
  - The Topic Operator and User Operator by configuring the **Kafka** custom resource:
    - Topic Operator
    - User Operator
- Alternative standalone deployment procedures for the Topic Operator and User Operator:
  - Deploy the standalone Topic Operator
  - Deploy the standalone User Operator

When installing Kafka, AMQ Streams also installs a ZooKeeper cluster and adds the necessary configuration to connect Kafka with ZooKeeper.

5.1.2.1. Deploying the Kafka cluster

This procedure shows how to deploy a Kafka cluster to your OpenShift using the Cluster Operator.

The deployment uses a YAML file to provide the specification to create a **Kafka** resource.

AMQ Streams provides example YAMLs files for deployment in **examples/kafka/**:

kafka-persistent.yaml

Deploys a persistent cluster with three ZooKeeper and three Kafka nodes.

kafka-jbod.yaml

Deploys a persistent cluster with three ZooKeeper and three Kafka nodes (each using multiple persistent volumes).

kafka-persistent-single.yaml

Deploys a persistent cluster with a single ZooKeeper node and a single Kafka node.

kafka-ephemeral.yaml

Deploys an ephemeral cluster with three ZooKeeper and three Kafka nodes.

kafka-ephemeral-single.yaml

Deploys an ephemeral cluster with three ZooKeeper nodes and a single Kafka node.

In this procedure, we use the examples for an *ephemeral* and *persistent* Kafka cluster deployment:

Ephemeral cluster

In general, an ephemeral (or temporary) Kafka cluster is suitable for development and testing purposes, not for production. This deployment uses **emptyDir** volumes for storing broker information (for ZooKeeper) and topics or partitions (for Kafka). Using an **emptyDir** volume means that its content is strictly related to the pod life cycle and is deleted when the pod goes down.

Persistent cluster

A persistent Kafka cluster uses **PersistentVolumes** to store ZooKeeper and Kafka data. The **PersistentVolume** is acquired using a **PersistentVolumeClaim** to make it independent of the actual type of the **PersistentVolume**. For example, it can use Amazon EBS volumes in Amazon AWS deployments without any changes in the YAML files. The **PersistentVolumeClaim** can use a **StorageClass** to trigger automatic volume provisioning.

The example YAML files specify the latest supported Kafka version, and configuration for its supported log message format version and inter-broker protocol version. Updates to these properties are required when upgrading Kafka.

The example clusters are named **my-cluster** by default. The cluster name is defined by the name of the resource and cannot be changed after the cluster has been deployed. To change the cluster name before you deploy the cluster, edit the **Kafka.metadata.name** property of the **Kafka** resource in the relevant YAML file.

Default cluster name and specified Kafka versions

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: my-cluster
```

```
spec:
  kafka:
    version: 2.7.0
    #...
    config:
      #...
      log.message.format.version: 2.7
      inter.broker.protocol.version: 2.7
  # ...
```

For more information about configuring the **Kafka** resource, see Kafka cluster configuration in the *Using AMQ Streams on OpenShift* guide.

Prerequisites

- The Cluster Operator must be deployed.

Procedure

1. Create and deploy an *ephemeral* or *persistent* cluster.
   For development or testing, you might prefer to use an ephemeral cluster. You can use a persistent cluster in any situation.
   - To create and deploy an *ephemeral* cluster:

     ```
     oc apply -f examples/kafka/kafka-ephemeral.yaml
     ```

   - To create and deploy a *persistent* cluster:

     ```
     oc apply -f examples/kafka/kafka-persistent.yaml
     ```

2. Verify that the Kafka cluster was successfully deployed:

   ```
   oc get deployments
   ```

5.1.2.2. Deploying the Topic Operator using the Cluster Operator

This procedure describes how to deploy the Topic Operator using the Cluster Operator.

You configure the **entityOperator** property of the **Kafka** resource to include the **topicOperator**.

If you want to use the Topic Operator with a Kafka cluster that is not managed by AMQ Streams, you must deploy the Topic Operator as a standalone component .

For more information about configuring the **entityOperator** and **topicOperator** properties, see Configuring the Entity Operator in the *Using AMQ Streams on OpenShift* guide.

Prerequisites

- The Cluster Operator must be deployed.

Procedure

1. Edit the **entityOperator** properties of the **Kafka** resource to include **topicOperator**:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: my-cluster
spec:
  #...
  entityOperator:
    topicOperator: {}
    userOperator: {}
```

2. Configure the Topic Operator **spec** using the properties described in **EntityTopicOperatorSpec** schema reference.
   Use an empty object (**{}**) if you want all properties to use their default values.

3. Create or update the resource:
   Use **oc apply**:

```
oc apply -f <your-file>
```

5.1.2.3. Deploying the User Operator using the Cluster Operator

This procedure describes how to deploy the User Operator using the Cluster Operator.

You configure the **entityOperator** property of the **Kafka** resource to include the **userOperator**.

If you want to use the User Operator with a Kafka cluster that is not managed by AMQ Streams, you must deploy the User Operator as a standalone component.

For more information about configuring the **entityOperator** and **userOperator** properties, see Configuring the Entity Operator in the *Using AMQ Streams on OpenShift* guide.

Prerequisites

- The Cluster Operator must be deployed.

Procedure

1. Edit the **entityOperator** properties of the **Kafka** resource to include **userOperator**:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: my-cluster
spec:
  #...
  entityOperator:
    topicOperator: {}
    userOperator: {}
```

2. Configure the User Operator **spec** using the properties described in **EntityUserOperatorSpec** schema reference in the *Using AMQ Streams on OpenShift* guide.
   Use an empty object (**{}**) if you want all properties to use their default values.

3. Create or update the resource:

```
oc apply -f <your-file>
```

5.1.3. Alternative standalone deployment options for AMQ Streams Operators

When deploying a Kafka cluster using the Cluster Operator, you can also deploy the Topic Operator and User Operator. Alternatively, you can perform a standalone deployment.

A standalone deployment means the Topic Operator and User Operator can operate with a Kafka cluster that is not managed by AMQ Streams.

5.1.3.1. Deploying the standalone Topic Operator

This procedure shows how to deploy the Topic Operator as a standalone component.

A standalone deployment requires configuration of environment variables, and is more complicated than deploying the Topic Operator using the Cluster Operator . However, a standalone deployment is more flexible as the Topic Operator can operate with *any* Kafka cluster, not necessarily one deployed by the Cluster Operator.

Prerequisites

- You need an existing Kafka cluster for the Topic Operator to connect to.

Procedure

1. Edit the **Deployment.spec.template.spec.containers[0].env** properties in the **install/topic-operator/05-Deployment-strimzi-topic-operator.yaml** file by setting:

   a. **STRIMZI_KAFKA_BOOTSTRAP_SERVERS** to list the bootstrap brokers in your Kafka cluster, given as a comma-separated list of ***hostname*:*port*** pairs.

   b. **STRIMZI_ZOOKEEPER_CONNECT** to list the ZooKeeper nodes, given as a comma-separated list of ***hostname*:*port*** pairs. This should be the same ZooKeeper cluster that your Kafka cluster is using.

   c. **STRIMZI_NAMESPACE** to the OpenShift namespace in which you want the operator to watch for **KafkaTopic** resources.

   d. **STRIMZI_RESOURCE_LABELS** to the label selector used to identify the **KafkaTopic** resources managed by the operator.

   e. **STRIMZI_FULL_RECONCILIATION_INTERVAL_MS** to specify the interval between periodic reconciliations, in milliseconds.

   f. **STRIMZI_TOPIC_METADATA_MAX_ATTEMPTS** to specify the number of attempts at getting topic metadata from Kafka. The time between each attempt is defined as an exponential back-off. Consider increasing this value when topic creation could take more time due to the number of partitions or replicas. Default **6**.

   g. **STRIMZI_ZOOKEEPER_SESSION_TIMEOUT_MS** to the ZooKeeper session timeout, in milliseconds. For example, **10000**. Default **20000** (20 seconds).

   h. **STRIMZI_TOPICS_PATH** to the Zookeeper node path where the Topic Operator stores its metadata. Default **/strimzi/topics**.

i. **STRIMZI_TLS_ENABLED** to enable TLS support for encrypting the communication with Kafka brokers. Default **true**.

j. **STRIMZI_TRUSTSTORE_LOCATION** to the path to the truststore containing certificates for enabling TLS based communication. Mandatory only if TLS is enabled through **STRIMZI_TLS_ENABLED**.

k. **STRIMZI_TRUSTSTORE_PASSWORD** to the password for accessing the truststore defined by **STRIMZI_TRUSTSTORE_LOCATION**. Mandatory only if TLS is enabled through **STRIMZI_TLS_ENABLED**.

l. **STRIMZI_KEYSTORE_LOCATION** to the path to the keystore containing private keys for enabling TLS based communication. Mandatory only if TLS is enabled through **STRIMZI_TLS_ENABLED**.

m. **STRIMZI_KEYSTORE_PASSWORD** to the password for accessing the keystore defined by **STRIMZI_KEYSTORE_LOCATION**. Mandatory only if TLS is enabled through **STRIMZI_TLS_ENABLED**.

n. **STRIMZI_LOG_LEVEL** to the level for printing logging messages. The value can be set to: **ERROR**, **WARNING**, **INFO**, **DEBUG**, and **TRACE**. Default **INFO**.

o. **STRIMZI_JAVA_OPTS** *(optional)* to the Java options used for the JVM running the Topic Operator. An example is **-Xmx=512M -Xms=256M**.

p. **STRIMZI_JAVA_SYSTEM_PROPERTIES** *(optional)* to list the **-D** options which are set to the Topic Operator. An example is **-Djavax.net.debug=verbose -DpropertyName=value**.

2. Deploy the Topic Operator:

```
oc create -f install/topic-operator
```

3. Verify that the Topic Operator has been deployed successfully:

```
oc describe deployment strimzi-topic-operator
```

The Topic Operator is deployed when the **Replicas:** entry shows **1 available**.

NOTE

You may experience a delay with the deployment if you have a slow connection to the OpenShift cluster and the images have not been downloaded before.

5.1.3.2. Deploying the standalone User Operator

This procedure shows how to deploy the User Operator as a standalone component.

A standalone deployment requires configuration of environment variables, and is more complicated than deploying the User Operator using the Cluster Operator . However, a standalone deployment is more flexible as the User Operator can operate with *any* Kafka cluster, not necessarily one deployed by the Cluster Operator.

Prerequisites

- You need an existing Kafka cluster for the User Operator to connect to.

Procedure

1. Edit the following **Deployment.spec.template.spec.containers[0].env** properties in the **install/user-operator/05-Deployment-strimzi-user-operator.yaml** file by setting:

   a. **STRIMZI_KAFKA_BOOTSTRAP_SERVERS** to list the Kafka brokers, given as a comma-separated list of ***hostname:port*** pairs.

   b. **STRIMZI_ZOOKEEPER_CONNECT** to list the ZooKeeper nodes, given as a comma-separated list of ***hostname:port*** pairs. This must be the same ZooKeeper cluster that your Kafka cluster is using. Connecting to ZooKeeper nodes with TLS encryption is not supported.

   c. **STRIMZI_NAMESPACE** to the OpenShift namespace in which you want the operator to watch for **KafkaUser** resources.

   d. **STRIMZI_LABELS** to the label selector used to identify the **KafkaUser** resources managed by the operator.

   e. **STRIMZI_FULL_RECONCILIATION_INTERVAL_MS** to specify the interval between periodic reconciliations, in milliseconds.

   f. **STRIMZI_ZOOKEEPER_SESSION_TIMEOUT_MS** to the ZooKeeper session timeout, in milliseconds. For example, **10000**. Default **20000** (20 seconds).

   g. **STRIMZI_CA_CERT_NAME** to point to an OpenShift **Secret** that contains the public key of the Certificate Authority for signing new user certificates for TLS client authentication. The **Secret** must contain the public key of the Certificate Authority under the key **ca.crt**.

   h. **STRIMZI_CA_KEY_NAME** to point to an OpenShift **Secret** that contains the private key of the Certificate Authority for signing new user certificates for TLS client authentication. The **Secret** must contain the private key of the Certificate Authority under the key **ca.key**.

   i. **STRIMZI_CLUSTER_CA_CERT_SECRET_NAME** to point to an OpenShift **Secret** containing the public key of the Certificate Authority used for signing Kafka brokers certificates for enabling TLS-based communication. The **Secret** must contain the public key of the Certificate Authority under the key **ca.crt**. This environment variable is optional and should be set only if the communication with the Kafka cluster is TLS based.

   j. **STRIMZI_EO_KEY_SECRET_NAME** to point to an OpenShift **Secret** containing the private key and related certificate for TLS client authentication against the Kafka cluster. The **Secret** must contain the keystore with the private key and certificate under the key **entity-operator.p12**, and the related password under the key **entity-operator.password**. This environment variable is optional and should be set only if TLS client authentication is needed when the communication with the Kafka cluster is TLS based.

   k. **STRIMZI_CA_VALIDITY** the validity period for the Certificate Authority. Default is **365** days.

   l. **STRIMZI_CA_RENEWAL** the renewal period for the Certificate Authority.

   m. **STRIMZI_LOG_LEVEL** to the level for printing logging messages. The value can be set to: **ERROR**, **WARNING**, **INFO**, **DEBUG**, and **TRACE**. Default **INFO**.

   n. **STRIMZI_GC_LOG_ENABLED** to enable garbage collection (GC) logging. Default **true**. Default is **30** days to initiate certificate renewal before the old certificates expire.

o. **STRIMZI_JAVA_OPTS** *(optional)* to the Java options used for the JVM running User Operator. An example is **-Xmx=512M -Xms=256M**.

p. **STRIMZI_JAVA_SYSTEM_PROPERTIES** *(optional)* to list the **-D** options which are set to the User Operator. An example is **-Djavax.net.debug=verbose -DpropertyName=value**.

2. Deploy the User Operator:

```
oc create -f install/user-operator
```

3. Verify that the User Operator has been deployed successfully:

```
oc describe deployment strimzi-user-operator
```

The User Operator is deployed when the **Replicas:** entry shows **1 available**.

NOTE

You may experience a delay with the deployment if you have a slow connection to the OpenShift cluster and the images have not been downloaded before.

5.2. DEPLOY KAFKA CONNECT

Kafka Connect is a tool for streaming data between Apache Kafka and external systems.

In AMQ Streams, Kafka Connect is deployed in distributed mode. Kafka Connect can also work in standalone mode, but this is not supported by AMQ Streams.

Using the concept of *connectors*, Kafka Connect provides a framework for moving large amounts of data into and out of your Kafka cluster while maintaining scalability and reliability.

Kafka Connect is typically used to integrate Kafka with external databases and storage and messaging systems.

The procedures in this section show how to:

- Deploy a Kafka Connect cluster using a **KafkaConnect** resource
- Run multiple Kafka Connect instances
- Create a Kafka Connect image containing the connectors you need to make your connection
- Create and manage connectors using a KafkaConnector resource or the Kafka Connect REST API
- Deploy a KafkaConnector resource to Kafka Connect
- Restart a Kafka connector by annotating a KafkaConnector resource
- Restart a Kafka connector task by annotating a KafkaConnector resource

NOTE

The term *connector* is used interchangeably to mean a connector instance running within a Kafka Connect cluster, or a connector class. In this guide, the term *connector* is used when the meaning is clear from the context.

5.2.1. Deploying Kafka Connect to your OpenShift cluster

This procedure shows how to deploy a Kafka Connect cluster to your OpenShift cluster using the Cluster Operator.

A Kafka Connect cluster is implemented as a **Deployment** with a configurable number of nodes (also called *workers*) that distribute the workload of connectors as *tasks* so that the message flow is highly scalable and reliable.

The deployment uses a YAML file to provide the specification to create a **KafkaConnect** resource.

In this procedure, we use the example file provided with AMQ Streams:

- **examples/connect/kafka-connect.yaml**

For information about configuring the **KafkaConnect** resource (or the **KafkaConnectS2I** resource with Source-to-Image (S2I) support), see Kafka Connect cluster configuration in the *Using AMQ Streams on OpenShift* guide.

Prerequisites

- The Cluster Operator must be deployed.
- Running Kafka cluster.

Procedure

1. Deploy Kafka Connect to your OpenShift cluster. For a Kafka cluster with 3 or more brokers, use the **examples/connect/kafka-connect.yaml** file. For a Kafka cluster with less than 3 brokers, use the **examples/connect/kafka-connect-single-node-kafka.yaml** file.

```
oc apply -f examples/connect/kafka-connect.yaml
```

2. Verify that Kafka Connect was successfully deployed:

```
oc get deployments
```

5.2.2. Kafka Connect configuration for multiple instances

If you are running multiple instances of Kafka Connect, you have to change the default configuration of the following **config** properties:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaConnect
metadata:
  name: my-connect
spec:
  # ...
```

```
  config:
    group.id: connect-cluster 1
    offset.storage.topic: connect-cluster-offsets 2
    config.storage.topic: connect-cluster-configs 3
    status.storage.topic: connect-cluster-status 4
    # ...
# ...
```

1. Kafka Connect cluster group that the instance belongs to.
2. Kafka topic that stores connector offsets.
3. Kafka topic that stores connector and task status configurations.
4. Kafka topic that stores connector and task status updates.

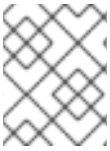

NOTE

Values for the three topics must be the same for all Kafka Connect instances with the same **group.id**.

Unless you change the default settings, each Kafka Connect instance connecting to the same Kafka cluster is deployed with the same values. What happens, in effect, is all instances are coupled to run in a cluster and use the same topics.

If multiple Kafka Connect clusters try to use the same topics, Kafka Connect will not work as expected and generate errors.

If you wish to run multiple Kafka Connect instances, change the values of these properties for each instance.

5.2.3. Extending Kafka Connect with connector plug-ins

The AMQ Streams container images for Kafka Connect include two built-in file connectors for moving file-based data into and out of your Kafka cluster.

Table 5.1. File connectors

File Connector	Description
FileStreamSourceConnector	Transfers data to your Kafka cluster from a file (the source).
FileStreamSinkConnector	Transfers data from your Kafka cluster to a file (the sink).

The procedures in this section show how to add your own connector classes to connector images by:

- Creating a new container image automatically using AMQ Streams
- Creating a container image from the Kafka Connect base image (manually or using continuous integration)

- Creating a container image using OpenShift builds and Source-to-Image (S2I) (available only on OpenShift)

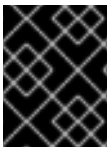

IMPORTANT

You create the configuration for connectors directly using the Kafka Connect REST API or KafkaConnector custom resources.

5.2.3.1. Creating a new container image automatically using AMQ Streams

This procedure shows how to configure Kafka Connect so that AMQ Streams automatically builds a new container image with additional connectors. You define the connector plugins using the **.spec.build.plugins** property of the **KafkaConnect** custom resource. AMQ Streams will automatically download and add the connector plugins into a new container image. The container is pushed into the container repository specified in **.spec.build.output** and automatically used in the Kafka Connect deployment.

Prerequisites

- The Cluster Operator must be deployed.
- A container registry.

You need to provide your own container registry where images can be pushed to, stored, and pulled from. AMQ Streams supports private container registries as well as public registries such as Quay or Docker Hub.

Procedure

1. Configure the **KafkaConnect** custom resource by specifying the container registry in **.spec.build.output**, and additional connectors in **.spec.build.plugins**:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaConnect
metadata:
  name: my-connect-cluster
spec: 1
  #...
  build:
    output: 2
      type: docker
      image: my-registry.io/my-org/my-connect-cluster:latest
      pushSecret: my-registry-credentials
    plugins: 3
      - name: debezium-postgres-connector
        artifacts:
          - type: tgz
            url: https://repo1.maven.org/maven2/io/debezium/debezium-connector-
postgres/1.3.1.Final/debezium-connector-postgres-1.3.1.Final-plugin.tar.gz
            sha512sum:
962a12151bdf9a5a30627eebac739955a4fd95a08d373b86bdcea2b4d0c27dd6e1edd5cb54804
5e115e33a9e69b1b2a352bee24df035a0447cb820077af00c03
      - name: camel-telegram
        artifacts:
          - type: tgz
```

```
        url: https://repo.maven.apache.org/maven2/org/apache/camel/kafkaconnector/camel-telegram-kafka-connector/0.7.0/camel-telegram-kafka-connector-0.7.0-package.tar.gz
        sha512sum: a9b1ac63e3284bea7836d7d24d84208c49cdf5600070e6bd1535de654f6920b74ad950d51733e8020bf4187870699819f54ef5859c7846ee4081507f48873479
      #...
```

1. The specification for the Kafka Connect cluster.
2. (Required) Configuration of the container registry where new images are pushed.
3. (Required) List of connector plugins and their artifacts to add to the new container image. Each plugin must be configured with at least one **artifact**.

2. Create or update the resource:

```
$ oc apply -f KAFKA-CONNECT-CONFIG-FILE
```

3. Wait for the new container image to build, and for the Kafka Connect cluster to be deployed.
4. Use the Kafka Connect REST API or the KafkaConnector custom resources to use the connector plugins you added.

Additional resources

See the *Using AMQ Streams on OpenShift* guide for more information on:

- Kafka Connect **Build** schema reference

5.2.3.2. Creating a Docker image from the Kafka Connect base image

This procedure shows how to create a custom image and add it to the **/opt/kafka/plugins** directory.

You can use the Kafka container image on Red Hat Ecosystem Catalog as a base image for creating your own custom image with additional connector plug-ins.

At startup, the AMQ Streams version of Kafka Connect loads any third-party connector plug-ins contained in the **/opt/kafka/plugins** directory.

Prerequisites

- The Cluster Operator must be deployed.

Procedure

1. Create a new **Dockerfile** using **registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0** as the base image:

```
FROM registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0
USER root:root
COPY ./my-plugins/ /opt/kafka/plugins/
USER 1001
```

Example plug-in file

```
$ tree ./my-plugins/
./my-plugins/
├── debezium-connector-mongodb
│   ├── bson-3.4.2.jar
│   ├── CHANGELOG.md
│   ├── CONTRIBUTE.md
│   ├── COPYRIGHT.txt
│   ├── debezium-connector-mongodb-0.7.1.jar
│   ├── debezium-core-0.7.1.jar
│   ├── LICENSE.txt
│   ├── mongodb-driver-3.4.2.jar
│   ├── mongodb-driver-core-3.4.2.jar
│   └── README.md
├── debezium-connector-mysql
│   ├── CHANGELOG.md
│   ├── CONTRIBUTE.md
│   ├── COPYRIGHT.txt
│   ├── debezium-connector-mysql-0.7.1.jar
│   ├── debezium-core-0.7.1.jar
│   ├── LICENSE.txt
│   ├── mysql-binlog-connector-java-0.13.0.jar
│   ├── mysql-connector-java-5.1.40.jar
│   ├── README.md
│   └── wkb-1.0.2.jar
└── debezium-connector-postgres
    ├── CHANGELOG.md
    ├── CONTRIBUTE.md
    ├── COPYRIGHT.txt
    ├── debezium-connector-postgres-0.7.1.jar
    ├── debezium-core-0.7.1.jar
    ├── LICENSE.txt
    ├── postgresql-42.0.0.jar
    ├── protobuf-java-2.6.1.jar
    └── README.md
```

2. Build the container image.

3. Push your custom image to your container registry.

4. Point to the new container image.
   You can either:

   - Edit the **KafkaConnect.spec.image** property of the **KafkaConnect** custom resource.
     If set, this property overrides the **STRIMZI_KAFKA_CONNECT_IMAGES** variable in the Cluster Operator.

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaConnect
metadata:
  name: my-connect-cluster
spec: 1
  #...
  image: my-new-container-image 2
  config: 3
    #...
```

1. The specification for the Kafka Connect cluster.
2. The docker image for the pods.
3. Configuration of the Kafka Connect *workers* (not connectors).

or

- In the **install/cluster-operator/060-Deployment-strimzi-cluster-operator.yaml** file, edit the **STRIMZI_KAFKA_CONNECT_IMAGES** variable to point to the new container image, and then reinstall the Cluster Operator.

Additional resources

See the *Using AMQ Streams on OpenShift* guide for more information on:

- Container image configuration and the **KafkaConnect.spec.image property**
- Cluster Operator configuration and the **STRIMZI_KAFKA_CONNECT_IMAGES** variable

5.2.3.3. Creating a container image using OpenShift builds and Source-to-Image

This procedure shows how to use OpenShift builds and the Source-to-Image (S2I) framework to create a new container image.

An OpenShift build takes a builder image with S2I support, together with source code and binaries provided by the user, and uses them to build a new container image. Once built, container images are stored in OpenShift's local container image repository and are available for use in deployments.

A Kafka Connect builder image with S2I support is provided on the Red Hat Ecosystem Catalog as part of the **registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0** image. This S2I image takes your binaries (with plug-ins and connectors) and stores them in the **/tmp/kafka-plugins/s2i** directory. It creates a new Kafka Connect image from this directory, which can then be used with the Kafka Connect deployment. When started using the enhanced image, Kafka Connect loads any third-party plug-ins from the **/tmp/kafka-plugins/s2i** directory.

IMPORTANT

With the introduction of **build** configuration to the **KafkaConnect** resource, AMQ Streams can now automatically build a container image with the connector plugins you require for your data connections. As a result, support for Kafka Connect with Source-to-Image (S2I) is deprecated. To prepare for this change, you can migrate Kafka Connect S2I instances to Kafka Connect instances.

Procedure

1. On the command line, use the **oc apply** command to create and deploy a Kafka Connect S2I cluster:

```
oc apply -f examples/connect/kafka-connect-s2i.yaml
```

2. Create a directory with Kafka Connect plug-ins:

```
$ tree ./my-plugins/
./my-plugins/
```

```
├── debezium-connector-mongodb
│   ├── bson-3.4.2.jar
│   ├── CHANGELOG.md
│   ├── CONTRIBUTE.md
│   ├── COPYRIGHT.txt
│   ├── debezium-connector-mongodb-0.7.1.jar
│   ├── debezium-core-0.7.1.jar
│   ├── LICENSE.txt
│   ├── mongodb-driver-3.4.2.jar
│   ├── mongodb-driver-core-3.4.2.jar
│   └── README.md
├── debezium-connector-mysql
│   ├── CHANGELOG.md
│   ├── CONTRIBUTE.md
│   ├── COPYRIGHT.txt
│   ├── debezium-connector-mysql-0.7.1.jar
│   ├── debezium-core-0.7.1.jar
│   ├── LICENSE.txt
│   ├── mysql-binlog-connector-java-0.13.0.jar
│   ├── mysql-connector-java-5.1.40.jar
│   ├── README.md
│   └── wkb-1.0.2.jar
└── debezium-connector-postgres
    ├── CHANGELOG.md
    ├── CONTRIBUTE.md
    ├── COPYRIGHT.txt
    ├── debezium-connector-postgres-0.7.1.jar
    ├── debezium-core-0.7.1.jar
    ├── LICENSE.txt
    ├── postgresql-42.0.0.jar
    ├── protobuf-java-2.6.1.jar
    └── README.md
```

3. Use the **oc start-build** command to start a new build of the image using the prepared directory:

```
oc start-build my-connect-cluster-connect --from-dir ./my-plugins/
```

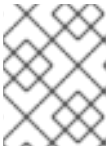

NOTE

The name of the build is the same as the name of the deployed Kafka Connect cluster.

4. When the build has finished, the new image is used automatically by the Kafka Connect deployment.

5.2.4. Creating and managing connectors

When you have created a container image for your connector plug-in, you need to create a connector instance in your Kafka Connect cluster. You can then configure, monitor, and manage a running connector instance.

A connector is an instance of a particular *connector class* that knows how to communicate with the relevant external system in terms of messages. Connectors are available for many external systems, or you can create your own.

You can create *source* and *sink* types of connector.

Source connector

A source connector is a runtime entity that fetches data from an external system and feeds it to Kafka as messages.

Sink connector

A sink connector is a runtime entity that fetches messages from Kafka topics and feeds them to an external system.

AMQ Streams provides two APIs for creating and managing connectors:

- KafkaConnector resources (referred to as KafkaConnectors)
- Kafka Connect REST API

Using the APIs, you can:

- Check the status of a connector instance
- Reconfigure a running connector
- Increase or decrease the number of connector tasks for a connector instance
- Restart connectors
- Restart connector tasks, including failed tasks
- Pause a connector instance
- Resume a previously paused connector instance
- Delete a connector instance

5.2.4.1. KafkaConnector resources

KafkaConnectors allow you to create and manage connector instances for Kafka Connect in an OpenShift-native way, so an HTTP client such as cURL is not required. Like other Kafka resources, you declare a connector's desired state in a KafkaConnector YAML file that is deployed to your OpenShift cluster to create the connector instance. KafkaConnector resources must be deployed to the same namespace as the Kafka Connect cluster they link to.

You manage a running connector instance by updating its corresponding KafkaConnector resource, and then applying the updates. Annotations are used to manually restart connector instances and connector tasks. You remove a connector by deleting its corresponding KafkaConnector.

To ensure compatibility with earlier versions of AMQ Streams, KafkaConnectors are disabled by default. To enable them for a Kafka Connect cluster, you must use annotations on the **KafkaConnect** resource. For instructions, see Configuring Kafka Connect in the *Using AMQ Streams on OpenShift* guide.

When KafkaConnectors are enabled, the Cluster Operator begins to watch for them. It updates the configurations of running connector instances to match the configurations defined in their KafkaConnectors.

AMQ Streams includes an example **KafkaConnector**, named **examples/connect/source-connector.yaml**. You can use this example to create and manage a **FileStreamSourceConnector** and a **FileStreamSinkConnector** as described in Section 5.2.5, "Deploying the example KafkaConnector

resources".

5.2.4.2. Availability of the Kafka Connect REST API

The Kafka Connect REST API is available on port 8083 as the **<connect-cluster-name>-connect-api** service.

If KafkaConnectors are enabled, manual changes made directly using the Kafka Connect REST API are reverted by the Cluster Operator.

The operations supported by the REST API are described in the Apache Kafka documentation.

5.2.5. Deploying the example KafkaConnector resources

AMQ Streams includes an example **KafkaConnector** in **examples/connect/source-connector.yaml**. This creates a basic **FileStreamSourceConnector** instance that sends each line of the Kafka license file (an example file source) to a single Kafka topic.

This procedure describes how to create:

- A **FileStreamSourceConnector** that reads data from the Kafka license file (the source) and writes the data as messages to a Kafka topic.
- A **FileStreamSinkConnector** that reads messages from the Kafka topic and writes the messages to a temporary file (the sink).

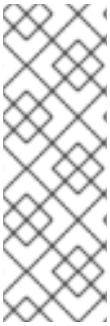

NOTE

In a production environment, you prepare container images containing your desired Kafka Connect connectors, as described in Section 5.2.3, "Extending Kafka Connect with connector plug-ins".

The **FileStreamSourceConnector** and **FileStreamSinkConnector** are provided as examples. Running these connectors in containers as described here is unlikely to be suitable for production use cases.

Prerequisites

- A Kafka Connect deployment
- KafkaConnectors are enabled in the Kafka Connect deployment
- The Cluster Operator is running

Procedure

1. Edit the **examples/connect/source-connector.yaml** file:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaConnector
metadata:
  name: my-source-connector 1
  labels:
    strimzi.io/cluster: my-connect-cluster 2
spec:
```

```
    class: org.apache.kafka.connect.file.FileStreamSourceConnector 3
    tasksMax: 2 4
    config: 5
      file: "/opt/kafka/LICENSE" 6
      topic: my-topic 7
      # ...
```

1. Name of the KafkaConnector resource, which is used as the name of the connector. Use any name that is valid for an OpenShift resource.
2. Name of the Kafka Connect cluster to create the connector instance in. Connectors must be deployed to the same namespace as the Kafka Connect cluster they link to.
3. Full name or alias of the connector class. This should be present in the image being used by the Kafka Connect cluster.
4. Maximum number of Kafka Connect **Tasks** that the connector can create.
5. Connector configuration as key-value pairs.
6. This example source connector configuration reads data from the **/opt/kafka/LICENSE** file.
7. Kafka topic to publish the source data to.

2. Create the source **KafkaConnector** in your OpenShift cluster:

```
oc apply -f examples/connect/source-connector.yaml
```

3. Create an **examples/connect/sink-connector.yaml** file:

```
touch examples/connect/sink-connector.yaml
```

4. Paste the following YAML into the **sink-connector.yaml** file:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaConnector
metadata:
  name: my-sink-connector
  labels:
    strimzi.io/cluster: my-connect
spec:
  class: org.apache.kafka.connect.file.FileStreamSinkConnector 1
  tasksMax: 2
  config: 2
    file: "/tmp/my-file" 3
    topics: my-topic 4
```

1. Full name or alias of the connector class. This should be present in the image being used by the Kafka Connect cluster.
2. Connector configuration as key-value pairs.

3. Temporary file to publish the source data to.
4. Kafka topic to read the source data from.

5. Create the sink **KafkaConnector** in your OpenShift cluster:

```
oc apply -f examples/connect/sink-connector.yaml
```

6. Check that the connector resources were created:

```
oc get kctr --selector strimzi.io/cluster=MY-CONNECT-CLUSTER -o name

my-source-connector
my-sink-connector
```

Replace *MY-CONNECT-CLUSTER* with your Kafka Connect cluster.

7. In the container, execute **kafka-console-consumer.sh** to read the messages that were written to the topic by the source connector:

```
oc exec MY-CLUSTER-kafka-0 -i -t -- bin/kafka-console-consumer.sh --bootstrap-server MY-CLUSTER-kafka-bootstrap.NAMESPACE.svc:9092 --topic my-topic --from-beginning
```

Source and sink connector configuration options
The connector configuration is defined in the **spec.config** property of the KafkaConnector resource.

The **FileStreamSourceConnector** and **FileStreamSinkConnector** classes support the same configuration options as the Kafka Connect REST API. Other connectors support different configuration options.

Table 5.2. Configuration options for the FileStreamSource connector class

Name	Type	Default value	Description
file	String	Null	Source file to write messages to. If not specified, the standard input is used.
topic	List	Null	The Kafka topic to publish data to.

Table 5.3. Configuration options for FileStreamSinkConnector class

Name	Type	Default value	Description
file	String	Null	Destination file to write messages to. If not specified, the standard output is used.

Name	Type	Default value	Description
topics	List	Null	One or more Kafka topics to read data from.
topics.regex	String	Null	A regular expression matching one or more Kafka topics to read data from.

Additional resources

- Section 5.2.4, "Creating and managing connectors"

5.2.6. Performing a restart of a Kafka connector

This procedure describes how to manually trigger a restart of a Kafka connector by using an OpenShift annotation.

Prerequisites

- The Cluster Operator is running.

Procedure

1. Find the name of the **KafkaConnector** custom resource that controls the Kafka connector you want to restart:

```
oc get KafkaConnector
```

2. To restart the connector, annotate the **KafkaConnector** resource in OpenShift. For example, using **oc annotate**:

```
oc annotate KafkaConnector KAFKACONNECTOR-NAME strimzi.io/restart=true
```

3. Wait for the next reconciliation to occur (every two minutes by default).
The Kafka connector is restarted, as long as the annotation was detected by the reconciliation process. When Kafka Connect accepts the restart request, the annotation is removed from the **KafkaConnector** custom resource.

Additional resources

- Creating and managing connectors in the *Deploying and Upgrading* guide.

5.2.7. Performing a restart of a Kafka connector task

This procedure describes how to manually trigger a restart of a Kafka connector task by using an OpenShift annotation.

Prerequisites

- The Cluster Operator is running.

Procedure

1. Find the name of the **KafkaConnector** custom resource that controls the Kafka connector task you want to restart:

```
oc get KafkaConnector
```

2. Find the ID of the task to be restarted from the **KafkaConnector** custom resource. Task IDs are non-negative integers, starting from 0.

```
oc describe KafkaConnector KAFKACONNECTOR-NAME
```

3. To restart the connector task, annotate the **KafkaConnector** resource in OpenShift. For example, using **oc annotate** to restart task 0:

```
oc annotate KafkaConnector KAFKACONNECTOR-NAME strimzi.io/restart-task=0
```

4. Wait for the next reconciliation to occur (every two minutes by default).
The Kafka connector task is restarted, as long as the annotation was detected by the reconciliation process. When Kafka Connect accepts the restart request, the annotation is removed from the **KafkaConnector** custom resource.

Additional resources

- Creating and managing connectors in the *Deploying and Upgrading* guide.

5.3. DEPLOY KAFKA MIRRORMAKER

The Cluster Operator deploys one or more Kafka MirrorMaker replicas to replicate data between Kafka clusters. This process is called mirroring to avoid confusion with the Kafka partitions replication concept. MirrorMaker consumes messages from the source cluster and republishes those messages to the target cluster.

5.3.1. Deploying Kafka MirrorMaker to your OpenShift cluster

This procedure shows how to deploy a Kafka MirrorMaker cluster to your OpenShift cluster using the Cluster Operator.

The deployment uses a YAML file to provide the specification to create a **KafkaMirrorMaker** or **KafkaMirrorMaker2** resource depending on the version of MirrorMaker deployed.

In this procedure, we use the example files provided with AMQ Streams:

- **examples/mirror-maker/kafka-mirror-maker.yaml**
- **examples/mirror-maker/kafka-mirror-maker-2.yaml**

For information about configuring **KafkaMirrorMaker** or **KafkaMirrorMaker2** resources, see Kafka MirrorMaker cluster configuration in the *Using AMQ Streams on OpenShift* guide.

Prerequisites

- The Cluster Operator must be deployed.

Procedure

1. Deploy Kafka MirrorMaker to your OpenShift cluster:
   For MirrorMaker:

   ```
   oc apply -f examples/mirror-maker/kafka-mirror-maker.yaml
   ```

   For MirrorMaker 2.0:

   ```
   oc apply -f examples/mirror-maker/kafka-mirror-maker-2.yaml
   ```

2. Verify that MirrorMaker was successfully deployed:

   ```
   oc get deployments
   ```

5.4. DEPLOY KAFKA BRIDGE

The Cluster Operator deploys one or more Kafka bridge replicas to send data between Kafka clusters and clients via HTTP API.

5.4.1. Deploying Kafka Bridge to your OpenShift cluster

This procedure shows how to deploy a Kafka Bridge cluster to your OpenShift cluster using the Cluster Operator.

The deployment uses a YAML file to provide the specification to create a **KafkaBridge** resource.

In this procedure, we use the example file provided with AMQ Streams:

- **examples/bridge/kafka-bridge.yaml**

For information about configuring the **KafkaBridge** resource, see Kafka Bridge cluster configuration in the *Using AMQ Streams on OpenShift* guide.

Prerequisites

- The Cluster Operator must be deployed.

Procedure

1. Deploy Kafka Bridge to your OpenShift cluster:

   ```
   oc apply -f examples/bridge/kafka-bridge.yaml
   ```

2. Verify that Kafka Bridge was successfully deployed:

   ```
   oc get deployments
   ```

CHAPTER 6. SETTING UP CLIENT ACCESS TO THE KAFKA CLUSTER

After you have deployed AMQ Streams, the procedures in this section explain how to:

- Deploy example producer and consumer clients, which you can use to verify your deployment
- Set up external client access to the Kafka cluster
  The steps to set up access to the Kafka cluster for a client outside OpenShift are more complex, and require familiarity with the Kafka component configuration procedures described in the *Using AMQ Streams on OpenShift* guide.

6.1. DEPLOYING EXAMPLE CLIENTS

This procedure shows how to deploy example producer and consumer clients that use the Kafka cluster you created to send and receive messages.

Prerequisites

- The Kafka cluster is available for the clients.

Procedure

1. Deploy a Kafka producer.

```
oc run kafka-producer -ti --image=registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0
--rm=true --restart=Never -- bin/kafka-console-producer.sh --broker-list cluster-name-kafka-
bootstrap:9092 --topic my-topic
```

2. Type a message into the console where the producer is running.
3. Press *Enter* to send the message.
4. Deploy a Kafka consumer.

```
oc run kafka-consumer -ti --image=registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0
--rm=true --restart=Never -- bin/kafka-console-consumer.sh --bootstrap-server cluster-name-
kafka-bootstrap:9092 --topic my-topic --from-beginning
```

5. Confirm that you see the incoming messages in the consumer console.

6.2. SETTING UP ACCESS FOR CLIENTS OUTSIDE OF OPENSHIFT

This procedure shows how to configure client access to a Kafka cluster from outside OpenShift.

Using the address of the Kafka cluster, you can provide external access to a client on a different OpenShift namespace or outside OpenShift entirely.

You configure an external Kafka listener to provide the access.

The following external listener types are supported:

- **route** to use OpenShift **Route** 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

The type chosen depends on your requirements, and your environment and infrastructure. For example, loadbalancers might not be suitable for certain infrastructure, such as bare metal, where node ports provide a better option.

In this procedure:

1. An external listener is configured for the Kafka cluster, with TLS encryption and authentication, and Kafka *simple authorization* is enabled.
2. A **KafkaUser** is created for the client, with TLS authentication and Access Control Lists (ACLs) defined for *simple authorization*.

You can configure your listener to use TLS or SCRAM-SHA-512 authentication, both of which can be used with TLS encryption. If you are using an authorization server, you can use token-based OAuth 2.0 authentication and OAuth 2.0 authorization. Open Policy Agent (OPA) authorization is also supported as a Kafka authorization option.

When you configure the **KafkaUser** authentication and authorization mechanisms, ensure they match the equivalent Kafka configuration:

- **KafkaUser.spec.authentication** matches **Kafka.spec.kafka.listeners[*].authentication**
- **KafkaUser.spec.authorization** matches **Kafka.spec.kafka.authorization**

You should have at least one listener supporting the authentication you want to use for the **KafkaUser**.

NOTE

Authentication between Kafka users and Kafka brokers depends on the authentication settings for each. For example, it is not possible to authenticate a user with TLS if it is not also enabled in the Kafka configuration.

AMQ Streams operators automate the configuration process:

- The Cluster Operator creates the listeners and sets up the cluster and client certificate authority (CA) certificates to enable authentication within the Kafka cluster.
- The User Operator creates the user representing the client and the security credentials used for client authentication, based on the chosen authentication type.

In this procedure, the certificates generated by the Cluster Operator are used, but you can replace them by installing your own certificates. You can also configure your listener to use a Kafka listener certificate managed by an external Certificate Authority.

Certificates are available in PKCS #12 format (.p12) and PEM (.crt) formats.

Prerequisites

- The Kafka cluster is available for the client

- The Cluster Operator and User Operator are running in the cluster
- A client outside the OpenShift cluster to connect to the Kafka cluster

Procedure

1. Configure the Kafka cluster with an **external** Kafka listener.
   - Define the authentication required to access the Kafka broker through the listener
   - Enable authorization on the Kafka broker
     For example:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: my-cluster
  namespace: myproject
spec:
  kafka:
    # ...
    listeners: 1
    - name: external 2
      port: 9094 3
      type: LISTENER-TYPE 4
      tls: true 5
      authentication:
        type: tls 6
      configuration:
        preferredNodePortAddressType: InternalDNS 7
        bootstrap and broker service overrides 8
        #...
    authorization: 9
      type: simple
      superUsers:
        - super-user-name 10
  # ...
```

1. Configuration options for enabling external listeners are described in the Generic Kafka listener schema reference.
2. Name to identify the listener. Must be unique within the Kafka cluster.
3. Port number used by the listener inside Kafka. The port number has to be unique within a given Kafka cluster. Allowed port numbers are 9092 and higher with the exception of ports 9404 and 9999, which are already used for Prometheus and JMX. Depending on the listener type, the port number might not be the same as the port number that connects Kafka clients.
4. External listener type specified as **route**, **loadbalancer**, **nodeport** or **ingress**. An internal listener is specified as **internal**.
5. Enables TLS encryption on the listener. Default is **false**. TLS encryption is not required for **route** listeners.

6 Authentication specified as **tls**.

7 (Optional, for **nodeport** listeners only) Configuration to specify a preference for the first address type used by AMQ Streams as the node address.

8 (Optional) AMQ Streams automatically determines the addresses to advertise to clients. The addresses are automatically assigned by OpenShift. You can override bootstrap and broker service addresses if the infrastructure on which you are running AMQ Streams does not provide the right address. Validation is not performed on the overrides. The override configuration differs according to the listener type. For example, you can override hosts for **route**, DNS names or IP addresses for **loadbalancer**, and node ports for **nodeport**.

9 Authoization specified as **simple**, which uses the **AclAuthorizer** Kafka plugin.

10 (Optional) Super users can access all brokers regardless of any access restrictions defined in ACLs.

WARNING

An OpenShift Route address comprises the name of the Kafka cluster, the name of the listener, and the name of the namespace it is created in. For example, **my-cluster-kafka-listener1-bootstrap-myproject** (*CLUSTER-NAME*-kafka-*LISTENER-NAME*-bootstrap-*NAMESPACE*). If you are using a **route** listener type, be careful that the whole length of the address does not exceed a maximum limit of 63 characters.

2. Create or update the **Kafka** resource.

```
oc apply -f KAFKA-CONFIG-FILE
```

The Kafka cluster is configured with a Kafka broker listener using TLS authentication.

A service is created for each Kafka broker pod.

A service is created to serve as the *bootstrap address* for connection to the Kafka cluster.

A service is also created as the *external bootstrap address* for external connection to the Kafka cluster using **nodeport** listeners.

The cluster CA certificate to verify the identity of the kafka brokers is also created with the same name as the **Kafka** resource.

3. Find the bootstrap address and port from the status of the **Kafka** resource.

```
oc get kafka KAFKA-CLUSTER-NAME -o jsonpath='{.status.listeners[?
(@.type=="external")].bootstrapServers}'
```

Use the bootstrap address in your Kafka client to connect to the Kafka cluster.

4. Extract the public cluster CA certificate and password from the generated ***KAFKA-CLUSTER-NAME*-cluster-ca-cert** Secret.

```
oc get secret KAFKA-CLUSTER-NAME-cluster-ca-cert -o jsonpath='{.data.ca\.p12}' | base64 -d > ca.p12
```

```
oc get secret KAFKA-CLUSTER-NAME-cluster-ca-cert -o jsonpath='{.data.ca\.password}' | base64 -d > ca.password
```

Use the certificate and password in your Kafka client to connect to the Kafka cluster with TLS encryption.

NOTE

Cluster CA certificates renew automatically by default. If you are using your own Kafka listener certificates, you will need to renew the certificates manually.

5. Create or modify a user representing the client that requires access to the Kafka cluster.
   - Specify the same authentication type as the **Kafka** listener.
   - Specify the authorization ACLs for simple authorization.
     For example:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaUser
metadata:
  name: my-user
  labels:
    strimzi.io/cluster: my-cluster (1)
spec:
  authentication:
    type: tls (2)
  authorization:
    type: simple
    acls: (3)
      - resource:
          type: topic
          name: my-topic
          patternType: literal
        operation: Read
      - resource:
          type: topic
          name: my-topic
          patternType: literal
        operation: Describe
      - resource:
          type: group
          name: my-group
          patternType: literal
        operation: Read
```

(1) The label must match the label of the Kafka cluster for the user to be created.

2 Authentication specified as **tls**.

3 Simple authorization requires an accompanying list of ACL rules to apply to the user. The rules define the operations allowed on Kafka resources based on the *username* (**my-user**).

6. Create or modify the **KafkaUser** resource.

```
oc apply -f USER-CONFIG-FILE
```

The user is created, as well as a Secret with the same name as the **KafkaUser** resource. The Secret contains a private and public key for TLS client authentication.

For example:

```
apiVersion: v1
kind: Secret
metadata:
  name: my-user
  labels:
    strimzi.io/kind: KafkaUser
    strimzi.io/cluster: my-cluster
type: Opaque
data:
  ca.crt: PUBLIC-KEY-OF-THE-CLIENT-CA
  user.crt: USER-CERTIFICATE-CONTAINING-PUBLIC-KEY-OF-USER
  user.key: PRIVATE-KEY-OF-USER
  user.p12: P12-ARCHIVE-FILE-STORING-CERTIFICATES-AND-KEYS
  user.password: PASSWORD-PROTECTING-P12-ARCHIVE
```

7. Configure your client to connect to the Kafka cluster with the properties required to make a secure connection to the Kafka cluster.

   a. Add the authentication details for the public cluster certificates:

```
security.protocol: SSL 1
ssl.truststore.location: PATH-TO/ssl/keys/truststore 2
ssl.truststore.password: CLUSTER-CA-CERT-PASSWORD 3
ssl.truststore.type=PKCS12 4
```

1 Enables TLS encryption (with or without TLS client authentication).

2 Specifies the truststore location where the certificates were imported.

3 Specifies the password for accessing the truststore. This property can be omitted if it is not needed by the truststore.

4 Identifies the truststore type.

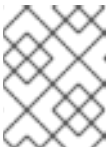

NOTE

Use **security.protocol: SASL_SSL** when using SCRAM-SHA authentication over TLS.

b. Add the bootstrap address and port for connecting to the Kafka cluster:

```
bootstrap.servers: BOOTSTRAP-ADDRESS:PORT
```

c. Add the authentication details for the public user certificates:

```
ssl.keystore.location: PATH-TO/ssl/keys/user1.keystore 1
ssl.keystore.password: USER-CERT-PASSWORD 2
```

1. Specifies the keystore location where the certificates were imported.
2. Specifies the password for accessing the keystore. This property can be omitted if it is not needed by the keystore.

The public user certificate is signed by the client CA when it is created.

CHAPTER 7. SETTING UP METRICS AND DASHBOARDS FOR AMQ STREAMS

You can monitor your AMQ Streams deployment by viewing key metrics on dashboards and setting up alerts that trigger under certain conditions. Metrics are available for Kafka, ZooKeeper, and the other components of AMQ Streams.

To provide metrics information, AMQ Streams uses Prometheus rules and Grafana dashboards.

When configured with a set of rules for each component of AMQ Streams, Prometheus consumes key metrics from the pods that are running in your cluster. Grafana then visualizes those metrics on dashboards. AMQ Streams includes example Grafana dashboards that you can customize to suit your deployment.

On OpenShift Container Platform 4.x, AMQ Streams employs *monitoring for user-defined projects* (an OpenShift feature) to simplify the Prometheus setup process.

On OpenShift Container Platform 3.11, you need to deploy the Prometheus and Alertmanager components to your cluster separately.

Regardless of your OpenShift Container Platform version, you have to start by deploying the Prometheus metrics configuration for AMQ Streams.

Next, follow the instructions for your OpenShift Container Platform version:

- Section 7.3, "Viewing Kafka metrics and dashboards in OpenShift 4"
- Section 7.4, "Viewing Kafka metrics and dashboards in OpenShift 3.11"

With Prometheus and Grafana set up, you can use the example Grafana dashboards and alerting rules to monitor your Kafka cluster.

Additional monitoring options

Kafka Exporter is an optional component that provides additional monitoring related to consumer lag. If you want to use Kafka Exporter with AMQ Streams, see Configure the **Kafka** resource to deploy Kafka Exporter with your Kafka cluster.

You can also configure your deployment to track messages end-to-end by setting up distributed tracing. For more information, see Distributed tracing in the *Using AMQ Streams on OpenShift* guide.

Additional resources

- Prometheus documentation
- Grafana documentation
- Apache Kafka Monitoring in the Kafka documentation describes JMX metrics exposed by Apache Kafka
- ZooKeeper JMX in the ZooKeeper documentation describes JMX metrics exposed by Apache ZooKeeper

7.1. EXAMPLE METRICS FILES

You can find example Grafana dashboards and other metrics configuration files in the **examples/metrics** directory. As indicated in the following list, some files are only used with OpenShift Container Platform 3.11, and not with OpenShift Container Platform 4.x.

Example metrics files provided with AMQ Streams

```
metrics
├── grafana-dashboards 1
│   ├── strimzi-cruise-control.json
│   ├── strimzi-kafka-bridge.json
│   ├── strimzi-kafka-connect.json
│   ├── strimzi-kafka-exporter.json
│   ├── strimzi-kafka-mirror-maker-2.json
│   ├── strimzi-kafka.json
│   ├── strimzi-operators.json
│   └── strimzi-zookeeper.json
├── grafana-install
│   └── grafana.yaml 2
├── prometheus-additional-properties
│   └── prometheus-additional.yaml - OPENSHIFT 3.11 ONLY 3
├── prometheus-alertmanager-config
│   └── alert-manager-config.yaml 4
├── prometheus-install
│   ├── alert-manager.yaml - OPENSHIFT 3.11 ONLY 5
│   ├── prometheus-rules.yaml 6
│   ├── prometheus.yaml - OPENSHIFT 3.11 ONLY 7
│   ├── strimzi-pod-monitor.yaml 8
├── kafka-bridge-metrics.yaml 9
├── kafka-connect-metrics.yaml 10
├── kafka-cruise-control-metrics.yaml 11
├── kafka-metrics.yaml 12
└── kafka-mirror-maker-2-metrics.yaml 13
```

1. Example Grafana dashboards.
2. Installation file for the Grafana image.
3. *OPENSHIFT 3.11 ONLY*: Additional Prometheus configuration to scrape metrics for CPU, memory, and disk volume usage, which comes directly from the OpenShift cAdvisor agent and kubelet on the nodes.
4. Hook definitions for sending notifications through Alertmanager.
5. *OPENSHIFT 3.11 ONLY*: Resources for deploying and configuring Alertmanager.
6. Alerting rules examples for use with Prometheus Alertmanager.
7. *OPENSHIFT 3.11 ONLY*: Installation resource file for the Prometheus image.
8. PodMonitor definitions translated by the Prometheus Operator into jobs for the Prometheus server to be able to scrape metrics data directly from pods.
9. Kafka Bridge resource with metrics enabled.

10 Metrics configuration that defines Prometheus JMX Exporter relabeling rules for Kafka Connect.

11 Metrics configuration that defines Prometheus JMX Exporter relabeling rules for Cruise Control.

12 Metrics configuration that defines Prometheus JMX Exporter relabeling rules for Kafka and ZooKeeper.

13 Metrics configuration that defines Prometheus JMX Exporter relabeling rules for Kafka Mirror Maker 2.0.

7.1.1. Example Grafana dashboards

Example Grafana dashboards are provided for monitoring the following resources:

AMQ Streams Kafka

Shows metrics for:

- Brokers online count
- Active controllers in the cluster count
- Unclean leader election rate
- Replicas that are online
- Under-replicated partitions count
- Partitions which are at their minimum in sync replica count
- Partitions which are under their minimum in sync replica count
- Partitions that do not have an active leader and are hence not writable or readable
- Kafka broker pods memory usage
- Aggregated Kafka broker pods CPU usage
- Kafka broker pods disk usage
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count
- Total incoming byte rate
- Total outgoing byte rate
- Incoming messages rate
- Total produce request rate
- Byte rate
- Produce request rate

- Fetch request rate
- Network processor average time idle percentage
- Request handler average time idle percentage
- Log size

AMQ Streams ZooKeeper

Shows metrics for:

- Quorum Size of Zookeeper ensemble
- Number of *alive* connections
- Queued requests in the server count
- Watchers count
- ZooKeeper pods memory usage
- Aggregated ZooKeeper pods CPU usage
- ZooKeeper pods disk usage
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count
- Amount of time it takes for the server to respond to a client request (maximum, minimum and average)

AMQ Streams Kafka Connect

Shows metrics for:

- Total incoming byte rate
- Total outgoing byte rate
- Disk usage
- JVM memory used
- JVM garbage collection time

AMQ Streams Kafka MirrorMaker 2

Shows metrics for:

- Number of connectors
- Number of tasks
- Total incoming byte rate

- Total outgoing byte rate
- Disk usage
- JVM memory used
- JVM garbage collection time

AMQ Streams Operators

Shows metrics for:

- Custom resources
- Successful custom resource reconciliations per hour
- Failed custom resource reconciliations per hour
- Reconciliations without locks per hour
- Reconciliations started hour
- Periodical reconciliations per hour
- Maximum reconciliation time
- Average reconciliation time
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count

Dashboards are also provided for the Kafka Bridge and Cruise Control components of AMQ Streams.

All the dashboards provide JVM metrics, as well as metrics that are specific to each component. For example, the Operators dashboard provides information on the number of reconciliations or custom resources that are being processed.

7.1.2. Example Prometheus metrics configuration

AMQ Streams uses the Prometheus JMX Exporter to expose JMX metrics using an HTTP endpoint, which is then scraped by Prometheus.

Grafana dashboards are dependent on Prometheus JMX Exporter relabeling rules, which are defined for AMQ Streams components as custom resource configuration.

A label is a name-value pair. Relabeling is the process of writing a label dynamically. For example, the value of a label might be derived from the name of a Kafka server and client ID.

AMQ Streams provides example custom resource configuration YAML files with the relabeling rules already defined. When deploying Prometheus metrics configuration, you can deploy the example custom resources or copy the metrics configuration to your own custom resource definitions.

Table 7.1. Example custom resources with metrics configuration

Component	Custom resource	Example YAML file
Kafka and ZooKeeper	**Kafka**	**kafka-metrics.yaml**
Kafka Connect	**KafkaConnect** and **KafkaConnectS2I**	**kafka-connect-metrics.yaml**
Kafka MirrorMaker 2.0	**KafkaMirrorMaker2**	**kafka-mirror-maker-2-metrics.yaml**
Kafka Bridge	**KafkaBridge**	**kafka-bridge-metrics.yaml**
Cruise Control	**Kafka**	**kafka-cruise-control-metrics.yaml**

Additional resources

- Section 7.2, "Deploying Prometheus metrics configuration"
- For more information on the use of relabeling, see Configuration in the Prometheus documentation.

7.2. DEPLOYING PROMETHEUS METRICS CONFIGURATION

AMQ Streams provides example custom resource configuration YAML files with relabeling rules.

To apply metrics configuration of relabeling rules, do one of the following:

- Copy the example configuration to your own custom resource definition
- Deploy the custom resource with the metrics configuration

7.2.1. Copying Prometheus metrics configuration to a custom resource

To use Grafana dashboards for monitoring, copy the example metrics configuration to a custom resource.

In this procedure, the **Kafka** resource is updated, but the procedure is the same for all components that support monitoring.

Procedure

Perform the following steps for each **Kafka** resource in your deployment.

1. Update the **Kafka** resource in an editor.

```
oc edit kafka KAFKA-CONFIG-FILE
```

2. Copy the example configuration in **kafka-metrics.yaml** to your own **Kafka** resource definition.
3. Save the file, and wait for the updated resource to be reconciled.

7.2.2. Deploying a Kafka cluster with Prometheus metrics configuration

To use Grafana dashboards for monitoring, you can deploy an example Kafka cluster with metrics configuration.

In this procedure, The **kafka-metrics.yaml** file is used for the **Kafka** resource.

Procedure

- Deploy the Kafka cluster with the example metrics configuration.

```
oc apply -f kafka-metrics.yaml
```

7.3. VIEWING KAFKA METRICS AND DASHBOARDS IN OPENSHIFT 4

When AMQ Streams is deployed to OpenShift Container Platform 4.x, metrics are provided through *monitoring for user-defined projects*. This OpenShift feature gives developers access to a separate Prometheus instance for monitoring their own projects (for example, a **Kafka** project).

If monitoring for user-defined projects is enabled, the **openshift-user-workload-monitoring** project contains the following components:

- A Prometheus Operator
- A Prometheus instance (automatically deployed by the Prometheus Operator)
- A Thanos Ruler instance

AMQ Streams uses these components to consume metrics.

A cluster administrator must enable monitoring for user-defined projects and then grant developers and other users permission to monitor applications within their own projects.

Grafana deployment

You can deploy a Grafana instance to the project containing your Kafka cluster. The example Grafana dashboards can then be used to visualize Prometheus metrics for AMQ Streams in the Grafana user interface.

IMPORTANT

The **openshift-monitoring** project provides monitoring for core platform components. Do *not* use the Prometheus and Grafana components in this project to configure monitoring for AMQ Streams on OpenShift Container Platform 4.x.

Grafana version 6.3 is the minimum supported version.

Prerequisites

- You have deployed the Prometheus metrics configuration using the example YAML files.
- *Monitoring for user-defined projects* is enabled. A cluster administrator must have created the **cluster-monitoring-config** ConfigMap in your OpenShift Container Platform cluster. For more information, see the following resources:
  - Enabling monitoring for user-defined projects in OpenShift Container Platform 4.6.

- Enabling monitoring of your own services in OpenShift Container Platform 4.5.

- To monitor user-defined projects, you must have been assigned the **monitoring-rules-edit** or **monitoring-edit** role by a cluster administrator. See:
  - Granting users permission to monitor user-defined projects in OpenShift Container Platform 4.6.
  - Granting user permissions using web console in OpenShift Container Platform 4.5.

Procedure outline

To set up AMQ Streams monitoring in OpenShift Container Platform 4.x, follow these procedures in order:

1. Prerequisite: Deploy the Prometheus metrics configuration
2. Deploy the Prometheus resources
3. Create a Service Account for Grafana
4. Deploy Grafana with a Prometheus datasource
5. Create a Route to the Grafana Service
6. Import the example Grafana dashboards

7.3.1. Deploying the Prometheus resources

NOTE

Use this procedure when running AMQ Streams on OpenShift Container Platform 4.x.

To enable Prometheus to consume Kafka metrics, you configure and deploy the **PodMonitor** resources in the example metrics files. The **PodMonitors** scrape data directly from pods for Apache Kafka, ZooKeeper, Operators, the Kafka Bridge, and Cruise Control.

Then, you deploy the example alerting rules for Alertmanager.

Prerequisites

- A running Kafka cluster.
- Check the example alerting rules provided with AMQ Streams.

Procedure

1. Check that monitoring for user-defined projects is enabled:

```
oc get pods -n openshift-user-workload-monitoring
```

If enabled, pods for the monitoring components are returned. For example:

```
NAME                                   READY   STATUS    RESTARTS   AGE
prometheus-operator-5cc59f9bc6-kgcq8   1/1     Running   0          25s
```

```
prometheus-user-workload-0        5/5     Running   1          14s
prometheus-user-workload-1        5/5     Running   1          14s
thanos-ruler-user-workload-0      3/3     Running   0          14s
thanos-ruler-user-workload-1      3/3     Running   0          14s
```

If no pods are returned, monitoring for user-defined projects is disabled. See the Prerequisites in Section 7.3, "Viewing Kafka metrics and dashboards in OpenShift 4" .

2. Multiple **PodMonitor** resources are defined in **examples/metrics/prometheus-install/strimzi-pod-monitor.yaml**.
For each **PodMonitor** resource, edit the **spec.namespaceSelector.matchNames** property:

```
apiVersion: monitoring.coreos.com/v1
kind: PodMonitor
metadata:
  name: cluster-operator-metrics
  labels:
    app: strimzi
spec:
  selector:
    matchLabels:
      strimzi.io/kind: cluster-operator
  namespaceSelector:
    matchNames:
      - PROJECT-NAME 1
  podMetricsEndpoints:
  - path: /metrics
    port: http
# ...
```

1 The project where the pods to scrape the metrics from are running, for example, **Kafka**.

3. Deploy the **strimzi-pod-monitor.yaml** file to the project where your Kafka cluster is running:

```
oc apply -f strimzi-pod-monitor.yaml -n MY-PROJECT
```

4. Deploy the example Prometheus rules to the same project:

```
oc apply -f prometheus-rules.yaml -n MY-PROJECT
```

Additional resources

- The Monitoring guide for OpenShift Container Platform 4.6
- Section 7.4.3.3, "Alerting rule examples"

7.3.2. Creating a Service Account for Grafana

NOTE

Use this procedure when running AMQ Streams on OpenShift Container Platform 4.x.

Your Grafana instance for AMQ Streams needs to run with a Service Account that is assigned the **cluster-monitoring-view** role.

Prerequisites

- Deploy the Prometheus resources

Procedure

1. Create a **ServiceAccount** for Grafana. Here the resource is named **grafana-serviceaccount**.

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: grafana-serviceaccount
  labels:
    app: strimzi
```

2. Deploy the **ServiceAccount** to the project containing your Kafka cluster:

```
oc apply -f GRAFANA-SERVICEACCOUNT -n MY-PROJECT
```

3. Create a **ClusterRoleBinding** resource that assigns the **cluster-monitoring-view** role to the Grafana **ServiceAccount**. Here the resource is named **grafana-cluster-monitoring-binding**.

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: grafana-cluster-monitoring-binding
  labels:
    app: strimzi
subjects:
  - kind: ServiceAccount
    name: grafana-serviceaccount
    namespace: MY-PROJECT (1)
roleRef:
  kind: ClusterRole
  name: cluster-monitoring-view
  apiGroup: rbac.authorization.k8s.io
```

(1) Name of your project.

4. Deploy the **ClusterRoleBinding** to the project containing your Kafka cluster:

```
oc apply -f GRAFANA-CLUSTER-MONITORING-BINDING -n MY-PROJECT
```

Additional resources

- Section 7.3, “Viewing Kafka metrics and dashboards in OpenShift 4”

7.3.3. Deploying Grafana with a Prometheus datasource

NOTE

Use this procedure when running AMQ Streams on OpenShift Container Platform 4.x.

This procedure describes how to deploy a Grafana application that is configured for the OpenShift Container Platform 4.x monitoring stack.

OpenShift Container Platform 4.x includes a *Thanos Querier* instance in the **openshift-monitoring** project. Thanos Querier is used to aggregate platform metrics.

To consume the required platform metrics, your Grafana instance requires a Prometheus data source that can connect to Thanos Querier. To configure this connection, you create a Config Map that authenticates, by using a token, to the **oauth-proxy** sidecar that runs alongside Thanos Querier. A **datasource.yaml** file is used as the source of the Config Map.

Finally, you deploy the Grafana application with the Config Map mounted as a volume to the project containing your Kafka cluster.

Prerequisites

- Deploy the Prometheus resources
- Create a Service Account for Grafana

Procedure

1. Get the access token of the Grafana **ServiceAccount**:

```
oc serviceaccounts get-token grafana-serviceaccount -n MY-PROJECT
```

   Copy the access token to use in the next step.

2. Create a **datasource.yaml** file containing the Thanos Querier configuration for Grafana. Paste the access token into the **httpHeaderValue1** property as indicated.

```
apiVersion: 1

datasources:
- name: Prometheus
  type: prometheus
  url: https://thanos-querier.openshift-monitoring.svc.cluster.local:9091
  access: proxy
  basicAuth: false
  withCredentials: false
  isDefault: true
  jsonData:
    timeInterval: 5s
    tlsSkipVerify: true
    httpHeaderName1: "Authorization"
  secureJsonData:
    httpHeaderValue1: "Bearer ${GRAFANA-ACCESS-TOKEN}" (1)
  editable: true
```

(1) ***GRAFANA-ACCESS-TOKEN***: The value of the access token for the Grafana **ServiceAccount**.

3. Create a Config Map named **grafana-config** from the **datasource.yaml** file:

```
oc create configmap grafana-config --from-file=datasource.yaml -n MY-PROJECT
```

4. Create a Grafana application consisting of a **Deployment** and a **Service**.
The **grafana-config** Config Map is mounted as a volume for the datasource configuration.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: grafana
  labels:
    app: strimzi
spec:
  replicas: 1
  selector:
    matchLabels:
      name: grafana
  template:
    metadata:
      labels:
        name: grafana
    spec:
      serviceAccountName: grafana-serviceaccount
      containers:
      - name: grafana
        image: grafana/grafana:6.3.0
        ports:
        - name: grafana
          containerPort: 3000
          protocol: TCP
        volumeMounts:
        - name: grafana-data
          mountPath: /var/lib/grafana
        - name: grafana-logs
          mountPath: /var/log/grafana
        - name: grafana-config
          mountPath: /etc/grafana/provisioning/datasources/datasource.yaml
          readOnly: true
          subPath: datasource.yaml
        readinessProbe:
          httpGet:
            path: /api/health
            port: 3000
          initialDelaySeconds: 5
          periodSeconds: 10
        livenessProbe:
          httpGet:
            path: /api/health
            port: 3000
          initialDelaySeconds: 15
          periodSeconds: 20
      volumes:
      - name: grafana-data
        emptyDir: {}
```

```
      - name: grafana-logs
        emptyDir: {}
      - name: grafana-config
        configMap:
          name: grafana-config
---
apiVersion: v1
kind: Service
metadata:
  name: grafana
  labels:
    app: strimzi
spec:
  ports:
  - name: grafana
    port: 3000
    targetPort: 3000
    protocol: TCP
  selector:
    name: grafana
  type: ClusterIP
```

5. Deploy the Grafana application to the project containing your Kafka cluster:

```
oc apply -f GRAFANA-APPLICATION -n MY-PROJECT
```

Additional resources

- Section 7.3, "Viewing Kafka metrics and dashboards in OpenShift 4"
- The Monitoring guide for OpenShift Container Platform 4.6

7.3.4. Creating a Route to the Grafana Service

NOTE

Use this procedure when running AMQ Streams on OpenShift Container Platform 4.x.

You can access the Grafana user interface through a Route that exposes the Grafana Service.

Prerequisites

- Deploy the Prometheus resources
- Create a Service Account for Grafana
- Deploy Grafana with a Prometheus datasource

Procedure

- Create an edge route to the **grafana** service:

```
oc create route edge MY-GRAFANA-ROUTE --service=grafana --namespace=KAFKA-NAMESPACE
```

Additional resources

- Section 7.3, "Viewing Kafka metrics and dashboards in OpenShift 4"

7.3.5. Importing the example Grafana dashboards

NOTE

Use this procedure when running AMQ Streams on OpenShift Container Platform 4.x.

Import the example Grafana dashboards using the Grafana user interface.

Prerequisites

- Deploy the Prometheus resources
- Create a Service Account for Grafana
- Deploy Grafana with a Prometheus datasource
- Create a Route to the Grafana Service

Procedure

1. Get the details of the Route to the Grafana Service. For example:

```
oc get routes

NAME              HOST/PORT                        PATH  SERVICES
MY-GRAFANA-ROUTE  MY-GRAFANA-ROUTE-amq-streams.net       grafana
```

2. In a web browser, access the Grafana login screen using the URL for the Route host and port.
3. Enter your user name and password, and then click **Log In**.
   The default Grafana user name and password are both **admin**. After logging in for the first time, you can change the password.
4. In **Configuration > Data Sources**, check that the **Prometheus** data source was created. The data source was created in Section 7.3.3, "Deploying Grafana with a Prometheus datasource".
5. Click **Dashboards > Manage**, and then click **Import**.
6. In **examples/metrics/grafana-dashboards**, copy the JSON of the dashboard to import.
7. Paste the JSON into the text box, and then click **Load**.
8. Repeat steps 1 -7 for the other example Grafana dashboards.

The imported Grafana dashboards are available to view from the **Dashboards** home page.

Additional resources

- Section 7.3.4, "Creating a Route to the Grafana Service"
- Section 7.3, "Viewing Kafka metrics and dashboards in OpenShift 4"

7.4. VIEWING KAFKA METRICS AND DASHBOARDS IN OPENSHIFT 3.11

When AMQ Streams is deployed to OpenShift Container Platform 3.11, you can use Prometheus to provide monitoring data for the example Grafana dashboards provided with AMQ Streams. You need to manually deploy the Prometheus components to your cluster.

In order to run the example Grafana dashboards, you must:

1. Add metrics configuration to your Kafka cluster resource
2. Deploy Prometheus and Prometheus Alertmanager
3. Deploy Grafana

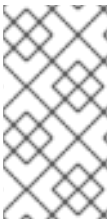

NOTE

The resources referenced in this section are intended as a starting point for setting up monitoring, but they are provided as examples only. If you require further support on configuring and running Prometheus or Grafana in production, try reaching out to their respective communities.

7.4.1. Prometheus support

The Prometheus server is *not* supported when AMQ Streams is deployed to OpenShift Container Platform 3.11. However, the Prometheus endpoint and the Prometheus JMX Exporter used to expose the metrics are supported.

For your convenience, we supply detailed instructions and example metrics configuration files should you wish to use Prometheus for monitoring.

7.4.2. Setting up Prometheus

NOTE

Use these procedures when running AMQ Streams on OpenShift Container Platform 3.11.

Prometheus provides an open source set of components for systems monitoring and alert notification.

Here we describe how to use the provided Prometheus image and configuration files to run and manage a Prometheus server when AMQ Streams is deployed to OpenShift Container Platform 3.11.

Prerequisites

- You have deployed compatible versions of Prometheus and Grafana to your OpenShift Container Platform 3.11 cluster.

- The service account used for running the Prometheus server pod has access to the OpenShift API server. This allows the service account to retrieve the list of pods in the cluster from which it gets metrics.
  For more information, see Discovering services.

7.4.2.1. Prometheus configuration

AMQ Streams provides example configuration files for the Prometheus server .

A Prometheus image is provided for deployment:

- **prometheus.yaml**

Additional Prometheus-related configuration is also provided in the following files:

- **prometheus-additional.yaml**
- **prometheus-rules.yaml**
- **strimzi-pod-monitor.yaml**

For Prometheus to obtain monitoring data, you must have deployed a compatible version of Prometheus to your OpenShift Container Platform 3.11 cluster.

Then, use the configuration files to Deploy Prometheus.

7.4.2.2. Prometheus resources

When you apply the Prometheus configuration, the following resources are created in your OpenShift cluster and managed by the Prometheus Operator:

- A **ClusterRole** that grants permissions to Prometheus to read the health endpoints exposed by the Kafka and ZooKeeper pods, cAdvisor and the kubelet for container metrics.
- A **ServiceAccount** for the Prometheus pods to run under.
- A **ClusterRoleBinding** which binds the **ClusterRole** to the **ServiceAccount**.
- A **Deployment** to manage the Prometheus Operator pod.
- A **PodMonitor** to manage the configuration of the Prometheus pod.
- A **Prometheus** to manage the configuration of the Prometheus pod.
- A **PrometheusRule** to manage alerting rules for the Prometheus pod.
- A **Secret** to manage additional Prometheus settings.
- A **Service** to allow applications running in the cluster to connect to Prometheus (for example, Grafana using Prometheus as datasource).

7.4.2.3. Deploying Prometheus

To obtain monitoring data in your Kafka cluster, you can use your own Prometheus deployment or deploy Prometheus by applying the example installation resource file for the Prometheus docker image and the YAML files for Prometheus-related resources.

The deployment process creates a **ClusterRoleBinding** and discovers an Alertmanager instance in the namespace specified for the deployment.

Prerequisites

- Check the example alerting rules provided

Procedure

1. Modify the Prometheus installation file (**prometheus.yaml**) according to the namespace Prometheus is going to be installed into:
   On Linux, use:

   ```
   sed -i 's/namespace: .*/namespace: my-namespace/' prometheus.yaml
   ```

   On MacOS, use:

   ```
   sed -i '' 's/namespace: .*/namespace: my-namespace/' prometheus.yaml
   ```

2. Edit the **PodMonitor** resource in **strimzi-pod-monitor.yaml** to define Prometheus jobs that will scrape the metrics data from pods.
   Update the **namespaceSelector.matchNames** property with the namespace where the pods to scrape the metrics from are running.

   PodMonitor is used to scrape data directly from pods for Apache Kafka, ZooKeeper, Operators, the Kafka Bridge and Cruise Control.

3. Edit the **prometheus.yaml** installation file to include additional configuration for scraping metrics directly from nodes.
   The Grafana dashboards provided show metrics for CPU, memory and disk volume usage, which come directly from the OpenShift cAdvisor agent and kubelet on the nodes.

   a. Create a **Secret** resource from the configuration file (**prometheus-additional.yaml** in the **examples/metrics/prometheus-additional-properties** directory):

      ```
      oc apply -f prometheus-additional.yaml
      ```

   b. Edit the **additionalScrapeConfigs** property in the **prometheus.yaml** file to include the name of the **Secret** and the **prometheus-additional.yaml** file.

4. Deploy the Prometheus resources:

   ```
   oc apply -f strimzi-pod-monitor.yaml
   oc apply -f prometheus-rules.yaml
   oc apply -f prometheus.yaml
   ```

7.4.3. Setting up Prometheus Alertmanager

Prometheus Alertmanager is a plugin for handling alerts and routing them to a notification service. Alertmanager supports an essential aspect of monitoring, which is to be notified of conditions that indicate potential issues based on alerting rules.

7.4.3.1. Alertmanager configuration

AMQ Streams provides example configuration files for Prometheus Alertmanager .

A configuration file defines the resources for deploying Alertmanager:

- **alert-manager.yaml**

An additional configuration file provides the hook definitions for sending notifications from your Kafka cluster.

- **alert-manager-config.yaml**

For Alertmanger to handle Prometheus alerts, use the configuration files to:

- Deploy Alertmanager

7.4.3.2. Alerting rules

Alerting rules provide notifications about specific conditions observed in the metrics. Rules are declared on the Prometheus server, but Prometheus Alertmanager is responsible for alert notifications.

Prometheus alerting rules describe conditions using PromQL expressions that are continuously evaluated.

When an alert expression becomes true, the condition is met and the Prometheus server sends alert data to the Alertmanager. Alertmanager then sends out a notification using the communication method configured for its deployment.

Alertmanager can be configured to use email, chat messages or other notification methods.

Additional resources

For more information about setting up alerting rules, see Configuration in the Prometheus documentation.

7.4.3.3. Alerting rule examples

Example alerting rules for Kafka and ZooKeeper metrics are provided with AMQ Streams for use in a Prometheus deployment.

General points about the alerting rule definitions:

- A **for** property is used with the rules to determine the period of time a condition must persist before an alert is triggered.
- A tick is a basic ZooKeeper time unit, which is measured in milliseconds and configured using the **tickTime** parameter of **Kafka.spec.zookeeper.config**. For example, if ZooKeeper **tickTime=3000**, 3 ticks (3 x 3000) equals 9000 milliseconds.
- The availability of the **ZookeeperRunningOutOfSpace** metric and alert is dependent on the OpenShift configuration and storage implementation used. Storage implementations for certain platforms may not be able to supply the information on available space required for the metric to provide an alert.

Kafka alerting rules

UnderReplicatedPartitions

Gives the number of partitions for which the current broker is the lead replica but which have fewer replicas than the **min.insync.replicas** configured for their topic. This metric provides insights about brokers that host the follower replicas. Those followers are not keeping up with the leader. Reasons for this could include being (or having been) offline, and over-throttled interbroker replication. An alert is raised when this value is greater than zero, providing information on the under-replicated partitions for each broker.

AbnormalControllerState

Indicates whether the current broker is the controller for the cluster. The metric can be 0 or 1. During the life of a cluster, only one broker should be the controller and the cluster always needs to have an active controller. Having two or more brokers saying that they are controllers indicates a problem. If the condition persists, an alert is raised when the sum of all the values for this metric on all brokers is not equal to 1, meaning that there is no active controller (the sum is 0) or more than one controller (the sum is greater than 1).

UnderMinIsrPartitionCount

Indicates that the minimum number of in-sync replicas (ISRs) for a lead Kafka broker, specified using **min.insync.replicas**, that must acknowledge a write operation has not been reached. The metric defines the number of partitions that the broker leads for which the in-sync replicas count is less than the minimum in-sync. An alert is raised when this value is greater than zero, providing information on the partition count for each broker that did not achieve the minimum number of acknowledgments.

OfflineLogDirectoryCount

Indicates the number of log directories which are offline (for example, due to a hardware failure) so that the broker cannot store incoming messages anymore. An alert is raised when this value is greater than zero, providing information on the number of offline log directories for each broker.

KafkaRunningOutOfSpace

Indicates the remaining amount of disk space that can be used for writing data. An alert is raised when this value is lower than 5GiB, providing information on the disk that is running out of space for each persistent volume claim. The threshold value may be changed in **prometheus-rules.yaml**.

ZooKeeper alerting rules

AvgRequestLatency

Indicates the amount of time it takes for the server to respond to a client request. An alert is raised when this value is greater than 10 (ticks), providing the actual value of the average request latency for each server.

OutstandingRequests

Indicates the number of queued requests in the server. This value goes up when the server receives more requests than it can process. An alert is raised when this value is greater than 10, providing the actual number of outstanding requests for each server.

ZookeeperRunningOutOfSpace

Indicates the remaining amount of disk space that can be used for writing data to ZooKeeper. An alert is raised when this value is lower than 5GiB., providing information on the disk that is running out of space for each persistent volume claim.

7.4.3.4. Deploying Alertmanager

To deploy Alertmanager, apply the example configuration files.

The sample configuration provided with AMQ Streams configures the Alertmanager to send notifications to a Slack channel.

The following resources are defined on deployment:

- An **Alertmanager** to manage the Alertmanager pod.
- A **Secret** to manage the configuration of the Alertmanager.
- A **Service** to provide an easy to reference hostname for other services to connect to Alertmanager (such as Prometheus).

Prerequisites

- Metrics are configured for the Kafka cluster resource
- Prometheus is deployed

Procedure

1. Create a **Secret** resource from the Alertmanager configuration file (**alert-manager-config.yaml** in the **examples/metrics/prometheus-alertmanager-config** directory):

```
oc apply -f alert-manager-config.yaml
```

2. Update the **alert-manager-config.yaml** file to replace the:
   - **slack_api_url** property with the actual value of the Slack API URL related to the application for the Slack workspace
   - **channel** property with the actual Slack channel on which to send notifications
3. Deploy Alertmanager:

```
oc apply -f alert-manager.yaml
```

7.4.4. Setting up Grafana

Grafana provides visualizations of Prometheus metrics.

You can deploy and enable the example Grafana dashboards provided with AMQ Streams.

7.4.4.1. Deploying Grafana

To provide visualizations of Prometheus metrics, you can use your own Grafana installation or deploy Grafana by applying the **grafana.yaml** file provided in the **examples/metrics** directory.

Prerequisites

- Metrics are configured for the Kafka cluster resource
- Prometheus and Prometheus Alertmanager are deployed

Procedure

1. Deploy Grafana:

```
oc apply -f grafana.yaml
```

2. Enable the Grafana dashboards.

7.4.4.2. Enabling the example Grafana dashboards

AMQ Streams provides example dashboard configuration files for Grafana . Example dashboards are provided in the **examples/metrics/grafana-dashboards** directory as JSON files:

- **strimzi-kafka.json**
- **strimzi-zookeeper.json**
- **strimzi-operators.json**
- **strimzi-kafka-connect.json**
- **strimzi-kafka-mirror-maker-2.json**
- **strimzi-kafka-bridge.json**
- **strimzi-cruise-control.json**
- **strimzi-kafka-exporter.json**

The example dashboards are a good starting point for monitoring key metrics, but they do not represent all available metrics. You can modify the example dashboards or add other metrics, depending on your infrastructure.

After setting up Prometheus and Grafana, you can visualize the AMQ Streams data on the Grafana dashboards.

NOTE

No alert notification rules are defined.

When accessing a dashboard, you can use the **port-forward** command to forward traffic from the Grafana pod to the host.

NOTE

The name of the Grafana pod is different for each user.

Procedure

1. Get the details of the Grafana service:

```
oc get service grafana
```

For example:

NAME	TYPE	CLUSTER-IP	PORT(S)
grafana	ClusterIP	172.30.123.40	3000/TCP

Note the port number for port forwarding.

2. Use **port-forward** to redirect the Grafana user interface to **localhost:3000**:

```
oc port-forward svc/grafana 3000:3000
```

3. Point a web browser to **http://localhost:3000**.
The Grafana Log In page appears.

4. Enter your user name and password, and then click **Log In**.
The default Grafana user name and password are both **admin**. After logging in for the first time, you can change the password.

5. Add Prometheus as a *data source*.

   - Specify a name
   - Add *Prometheus* as the type
   - Specify a Prometheus server URL (http://prometheus-operated:9090)
   Save and test the connection when you have added the details.

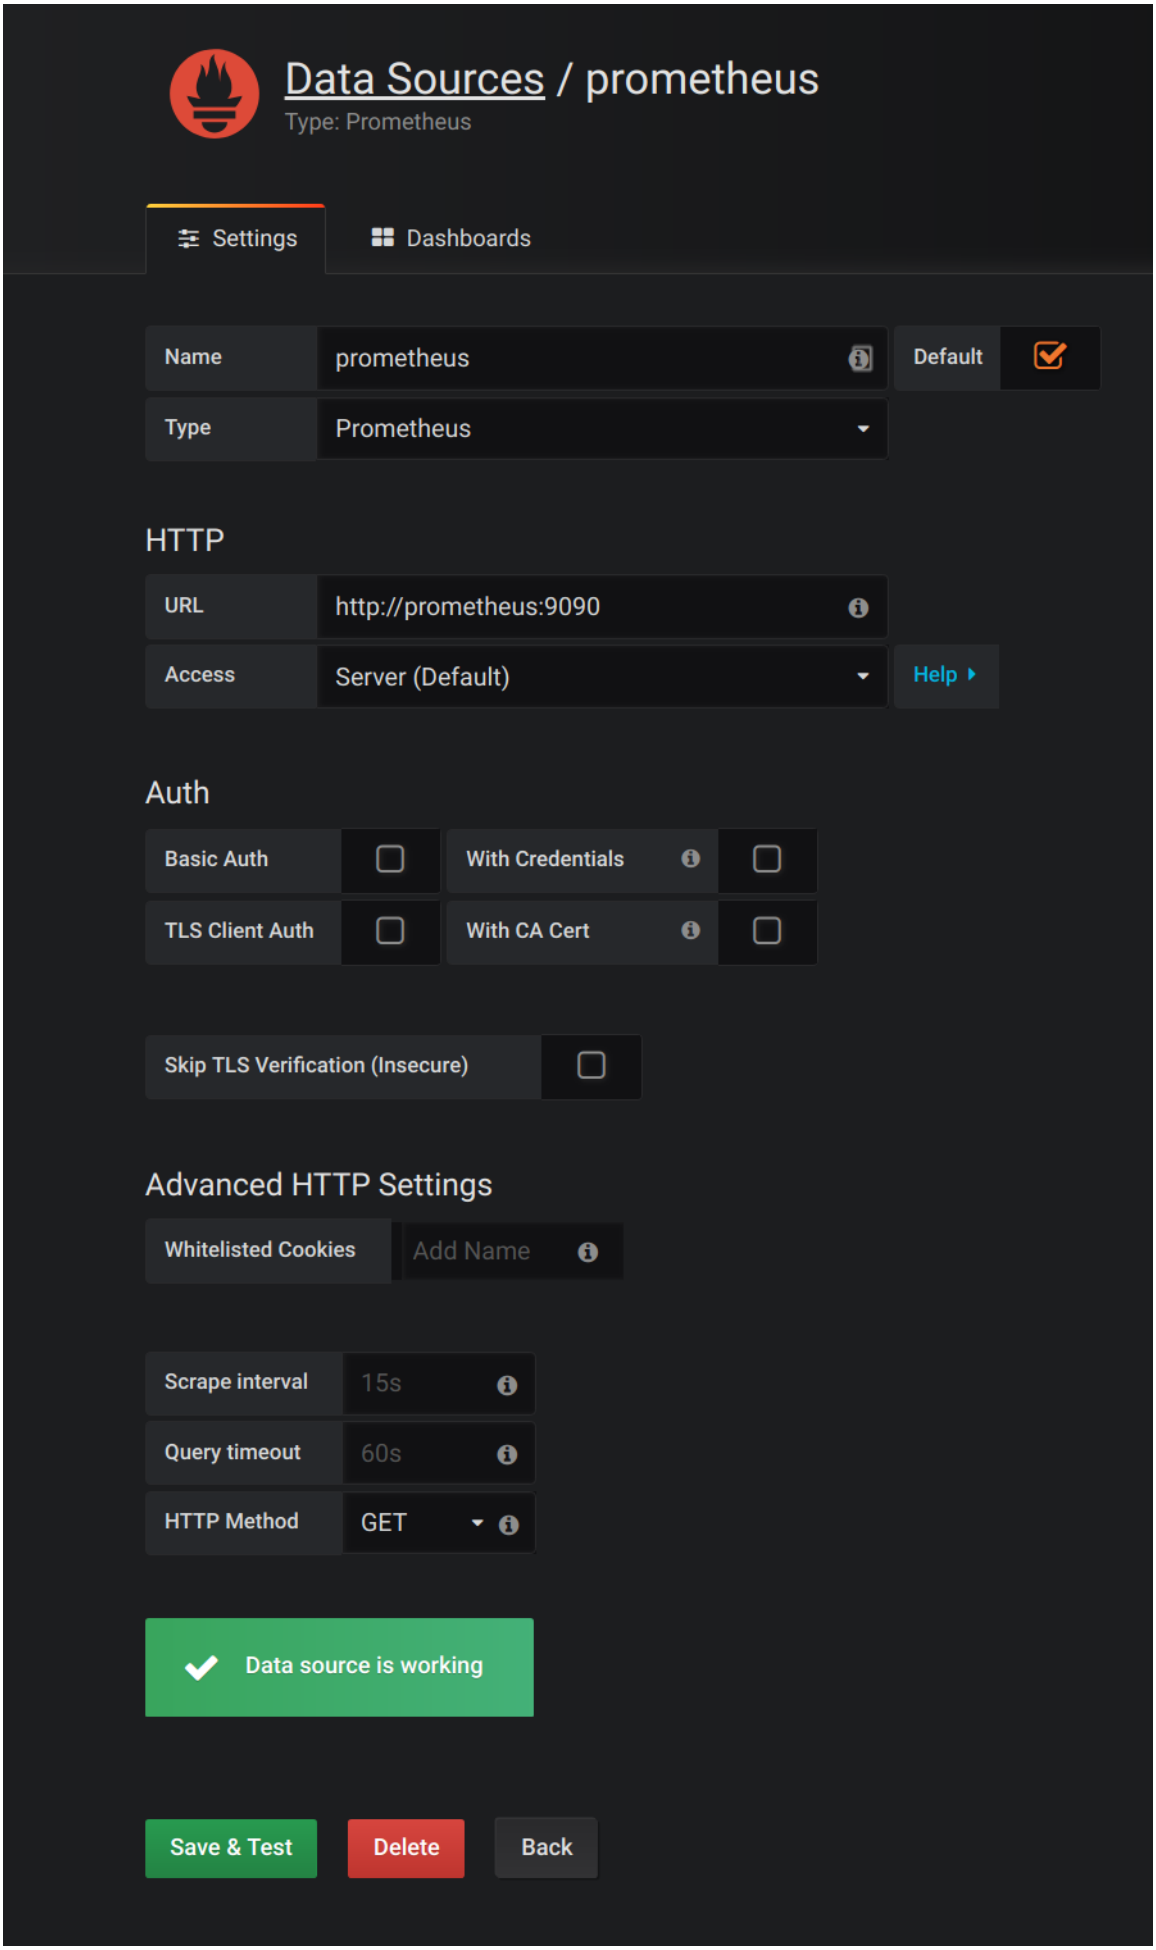

6. From **Dashboards → Import**, upload the example dashboards or paste the JSON directly.

7. On the top header, click the dashboard drop-down menu, and then select the dashboard you want to view.
   When the Prometheus server has been collecting metrics for a AMQ Streams cluster for some time, the dashboards are populated.

Figure 7.1. Dashboard selection options

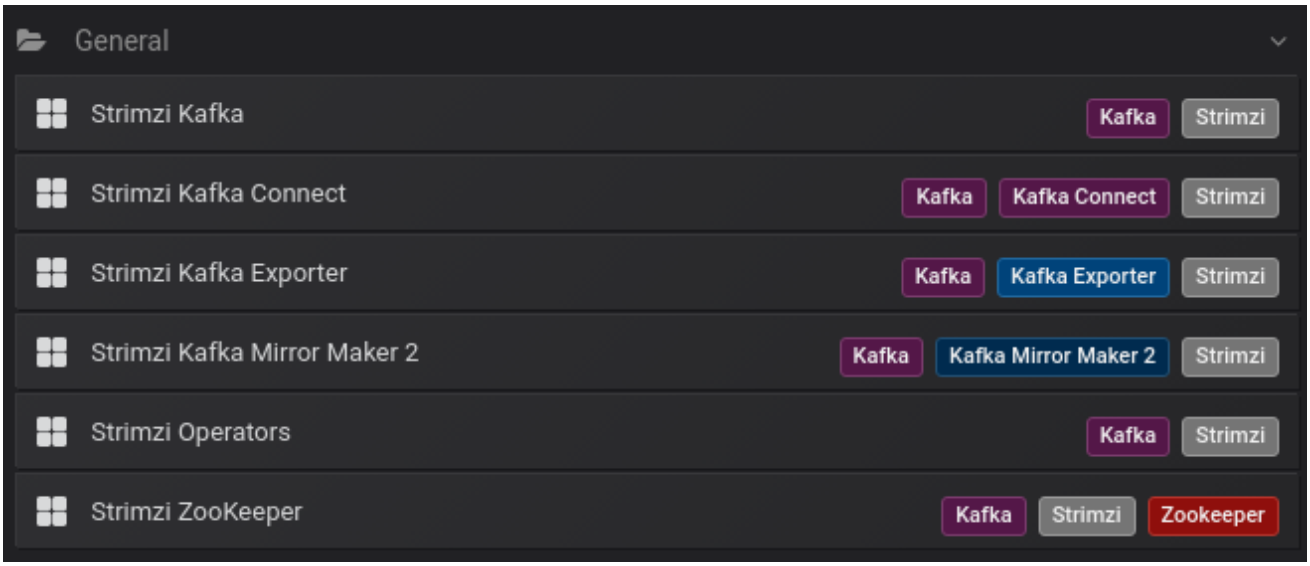

AMQ Streams Kafka

Shows metrics for:

- Brokers online count
- Active controllers in the cluster count
- Unclean leader election rate
- Replicas that are online
- Under-replicated partitions count
- Partitions which are at their minimum in sync replica count
- Partitions which are under their minimum in sync replica count
- Partitions that do not have an active leader and are hence not writable or readable
- Kafka broker pods memory usage
- Aggregated Kafka broker pods CPU usage
- Kafka broker pods disk usage
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count
- Total incoming byte rate
- Total outgoing byte rate

- Incoming messages rate
- Total produce request rate
- Byte rate
- Produce request rate
- Fetch request rate
- Network processor average time idle percentage
- Request handler average time idle percentage
- Log size

Figure 7.2. AMQ Streams Kafka dashboard

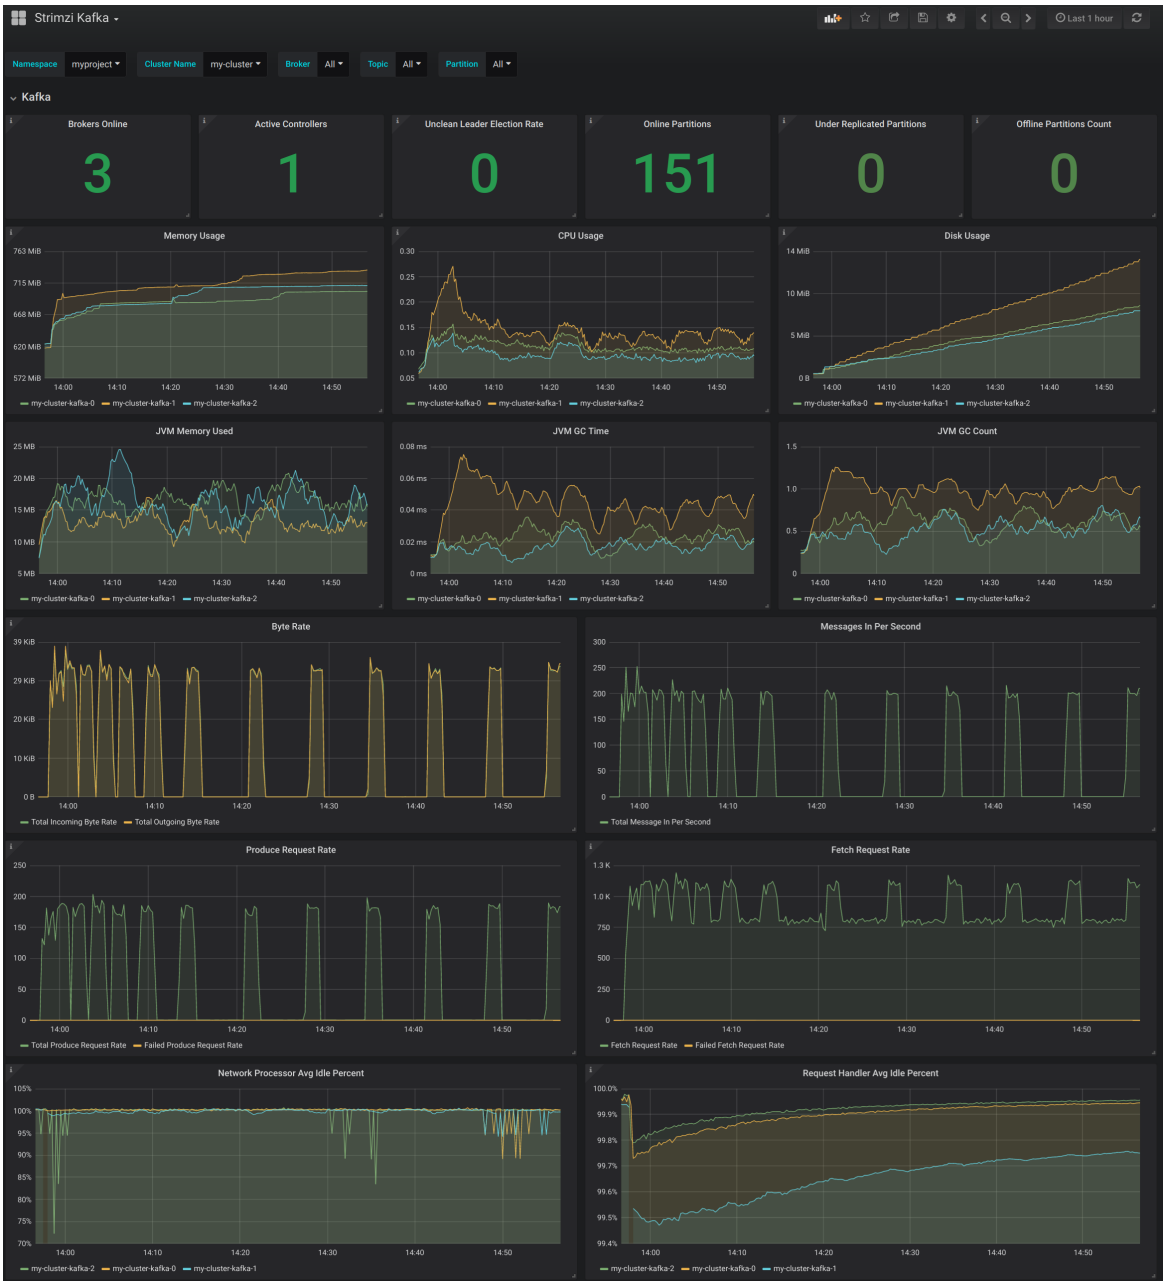

AMQ Streams ZooKeeper

Shows metrics for:

- Quorum Size of Zookeeper ensemble
- Number of *alive* connections
- Queued requests in the server count
- Watchers count
- ZooKeeper pods memory usage
- Aggregated ZooKeeper pods CPU usage
- ZooKeeper pods disk usage
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count
- Amount of time it takes for the server to respond to a client request (maximum, minimum and average)

AMQ Streams Operators

Shows metrics for:

- Custom resources
- Successful custom resource reconciliations per hour
- Failed custom resource reconciliations per hour
- Reconciliations without locks per hour
- Reconciliations started hour
- Periodical reconciliations per hour
- Maximum reconciliation time
- Average reconciliation time
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count

AMQ Streams Kafka Connect

Shows metrics for:

- Total incoming byte rate
- Total outgoing byte rate
- Disk usage

- JVM memory used
- JVM garbage collection time

AMQ Streams Kafka MirrorMaker 2

Shows metrics for:

- Number of connectors
- Number of tasks
- Total incoming byte rate
- Total outgoing byte rate
- Disk usage
- JVM memory used
- JVM garbage collection time

AMQ Streams Kafka Bridge

See Section 7.6, "Monitor Kafka Bridge" .

AMQ Streams Cruise Control

See Section 7.7, "Monitor Cruise Control".

AMQ Streams Kafka Exporter

See Section 7.5.5, "Enabling the Kafka Exporter Grafana dashboard" .

7.5. ADD KAFKA EXPORTER

Kafka Exporter is an open source project to enhance monitoring of Apache Kafka brokers and clients. Kafka Exporter is provided with AMQ Streams for deployment with a Kafka cluster to extract additional metrics data from Kafka brokers related to offsets, consumer groups, consumer lag, and topics.

The metrics data is used, for example, to help identify slow consumers.

Lag data is exposed as Prometheus metrics, which can then be presented in Grafana for analysis.

If you are already using Prometheus and Grafana for monitoring of built-in Kafka metrics, you can configure Prometheus to also scrape the Kafka Exporter Prometheus endpoint.

AMQ Streams includes an example Kafka Exporter dashboard in **examples/metrics/grafana-dashboards/strimzi-kafka-exporter.json**.

7.5.1. Monitoring Consumer lag

Consumer lag indicates the difference in the rate of production and consumption of messages. Specifically, consumer lag for a given consumer group indicates the delay between the last message in the partition and the message being currently picked up by that consumer.

The lag reflects the position of the consumer offset in relation to the end of the partition log.

Consumer lag between the producer and consumer offset

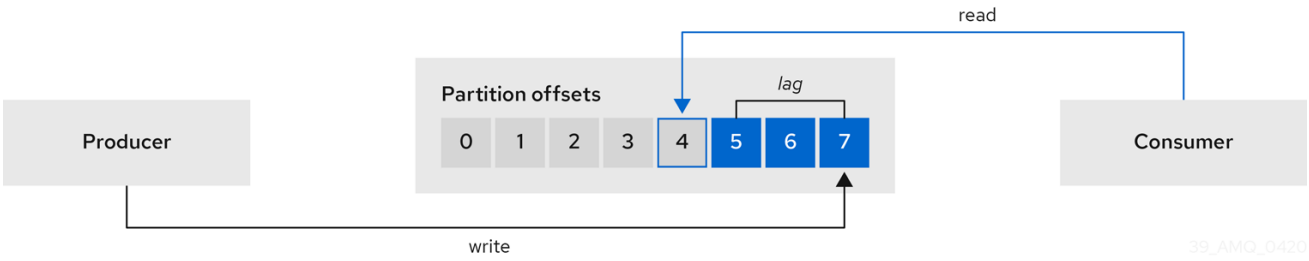

This difference is sometimes referred to as the *delta* between the producer offset and consumer offset: the read and write positions in the Kafka broker topic partitions.

Suppose a topic streams 100 messages a second. A lag of 1000 messages between the producer offset (the topic partition head) and the last offset the consumer has read means a 10-second delay.

The importance of monitoring consumer lag
For applications that rely on the processing of (near) real-time data, it is critical to monitor consumer lag to check that it does not become too big. The greater the lag becomes, the further the process moves from the real-time processing objective.

Consumer lag, for example, might be a result of consuming too much old data that has not been purged, or through unplanned shutdowns.

Reducing consumer lag
Typical actions to reduce lag include:

- Scaling-up consumer groups by adding new consumers
- Increasing the retention time for a message to remain in a topic
- Adding more disk capacity to increase the message buffer

Actions to reduce consumer lag depend on the underlying infrastructure and the use cases AMQ Streams is supporting. For instance, a lagging consumer is less likely to benefit from the broker being able to service a fetch request from its disk cache. And in certain cases, it might be acceptable to automatically drop messages until a consumer has caught up.

7.5.2. Example Kafka Exporter alerting rules

If you performed the steps to introduce metrics to your deployment, you will already have your Kafka cluster configured to use the alert notification rules that support Kafka Exporter.

The rules for Kafka Exporter are defined in **prometheus-rules.yaml**, and are deployed with Prometheus. For more information, see Prometheus.

The sample alert notification rules specific to Kafka Exporter are as follows:

UnderReplicatedPartition

An alert to warn that a topic is under-replicated and the broker is not replicating to enough partitions. The default configuration is for an alert if there are one or more under-replicated partitions for a topic. The alert might signify that a Kafka instance is down or the Kafka cluster is overloaded. A planned restart of the Kafka broker may be required to restart the replication process.

TooLargeConsumerGroupLag

An alert to warn that the lag on a consumer group is too large for a specific topic partition. The default configuration is 1000 records. A large lag might indicate that consumers are too slow and are falling behind the producers.

NoMessageForTooLong

An alert to warn that a topic has not received messages for a period of time. The default configuration for the time period is 10 minutes. The delay might be a result of a configuration issue preventing a producer from publishing messages to the topic.

Adapt the default configuration of these rules according to your specific needs.

Additional resources

- Chapter 7, *Setting up metrics and dashboards for AMQ Streams*
- Section 7.1, "Example metrics files"
- Section 7.4.3.2, "Alerting rules"

7.5.3. Exposing Kafka Exporter metrics

Lag information is exposed by Kafka Exporter as Prometheus metrics for presentation in Grafana.

Kafka Exporter exposes metrics data for brokers, topics and consumer groups. These metrics are displayed on the example **strimzi-kafka-exporter** dashboard.

The data extracted is described here.

Table 7.2. Broker metrics output

Name	Information
kafka_brokers	Number of brokers in the Kafka cluster

Table 7.3. Topic metrics output

Name	Information
kafka_topic_partitions	Number of partitions for a topic
kafka_topic_partition_current_offset	Current topic partition offset for a broker
kafka_topic_partition_oldest_offset	Oldest topic partition offset for a broker
kafka_topic_partition_in_sync_replica	Number of in-sync replicas for a topic partition
kafka_topic_partition_leader	Leader broker ID of a topic partition
kafka_topic_partition_leader_is_preferred	Shows **1** if a topic partition is using the preferred broker
kafka_topic_partition_replicas	Number of replicas for this topic partition

Name	Information
kafka_topic_partition_under_replicated_partition	Shows **1** if a topic partition is under-replicated

Table 7.4. Consumer group metrics output

Name	Information
kafka_consumergroup_current_offset	Current topic partition offset for a consumer group
kafka_consumergroup_lag	Current approximate lag for a consumer group at a topic partition

Consumer group metrics are only displayed on the Kafka Exporter dashboard if at least one consumer group has a lag greater than zero.

7.5.4. Configuring Kafka Exporter

This procedure shows how to configure Kafka Exporter in the **Kafka** resource through **KafkaExporter** properties.

For more information about configuring the **Kafka** resource, see Kafka cluster configuration in the *Using AMQ Streams on OpenShift* guide.

The properties relevant to the Kafka Exporter configuration are shown in this procedure.

You can configure these properties as part of a deployment or redeployment of the Kafka cluster.

Prerequisites

- An OpenShift cluster
- A running Cluster Operator

Procedure

1. Edit the **KafkaExporter** properties for the **Kafka** resource.
   The properties you can configure are shown in this example configuration:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: my-cluster
spec:
  # ...
  kafkaExporter:
    image: my-registry.io/my-org/my-exporter-cluster:latest 1
    groupRegex: ".*" 2
    topicRegex: ".*" 3
    resources: 4
```

```
      requests:
        cpu: 200m
        memory: 64Mi
      limits:
        cpu: 500m
        memory: 128Mi
    logging: debug 5
    enableSaramaLogging: true 6
    template: 7
      pod:
        metadata:
          labels:
            label1: value1
        imagePullSecrets:
          - name: my-docker-credentials
        securityContext:
          runAsUser: 1000001
          fsGroup: 0
        terminationGracePeriodSeconds: 120
    readinessProbe: 8
      initialDelaySeconds: 15
      timeoutSeconds: 5
    livenessProbe: 9
      initialDelaySeconds: 15
      timeoutSeconds: 5
# ...
```

1. ADVANCED OPTION: Container image configuration, which is recommended only in special situations.
2. A regular expression to specify the consumer groups to include in the metrics.
3. A regular expression to specify the topics to include in the metrics.
4. CPU and memory resources to reserve .
5. Logging configuration, to log messages with a given severity (debug, info, warn, error, fatal) or above.
6. Boolean to enable Sarama logging, a Go client library used by Kafka Exporter.
7. Customization of deployment templates and pods.
8. Healthcheck readiness probes.
9. Healthcheck liveness probes.

2. Create or update the resource:

```
oc apply -f kafka.yaml
```

What to do next

After configuring and deploying Kafka Exporter, you can enable Grafana to present the Kafka Exporter dashboards.

Additional resources

KafkaExporterTemplate schema reference.

7.5.5. Enabling the Kafka Exporter Grafana dashboard

AMQ Streams provides example dashboard configuration files for Grafana . The Kafka Exporter dashboard is provided in the **examples/metrics** directory as a JSON file:

- **strimzi-kafka-exporter.json**

If you deployed Kafka Exporter with your Kafka cluster, you can visualize the metrics data it exposes on the Grafana dashboard.

Prerequisites

- Kafka is deployed with Kafka Exporter metrics configuration
- Prometheus and Prometheus Alertmanager are deployed to the Kafka cluster
- Grafana is deployed to the Kafka cluster

This procedure assumes you already have access to the Grafana user interface and Prometheus has been added as a data source. If you are accessing the user interface for the first time, see Grafana.

Procedure

1. Access the Grafana user interface.
2. Select the *Strimzi Kafka Exporter* dashboard.
   When metrics data has been collected for some time, the Kafka Exporter charts are populated.

   AMQ Streams Kafka Exporter

   Shows metrics for:

   - Topic count
   - Partition count
   - Replicas count
   - In-sync replicas count
   - Under-replicated partitions count
   - Partitions which are at their minimum in sync replica count
   - Partitions which are under their minimum in sync replica count
   - Partitions not on a preferred node
   - Messages in per second from topics
   - Messages consumed per second from topics
   - Messages consumed per minute by consumer groups

- Lag by consumer group
- Number of partitions
- Latest offsets
- Oldest offsets

Use the Grafana charts to analyze lag and to check if actions to reduce lag are having an impact on an affected consumer group. If, for example, Kafka brokers are adjusted to reduce lag, the dashboard will show the *Lag by consumer group* chart going down and the *Messages consumed per minute* chart going up.

7.6. MONITOR KAFKA BRIDGE

If you are already using Prometheus and Grafana for monitoring of built-in Kafka metrics, you can configure Prometheus to also scrape the Kafka Bridge Prometheus endpoint.

The example Grafana dashboard for the Kafka Bridge provides:

- Information about HTTP connections and related requests to the different endpoints
- Information about the Kafka consumers and producers used by the bridge
- JVM metrics from the bridge itself

7.6.1. Configuring Kafka Bridge

You can enable the Kafka Bridge metrics in the **KafkaBridge** resource using the **enableMetrics** property.

You can configure this property as part of a deployment or redeployment of the Kafka Bridge.

For example:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: KafkaBridge
metadata:
  name: my-bridge
spec:
  # ...
  bootstrapServers: my-cluster-kafka:9092
  http:
    # ...
  enableMetrics: true
  # ...
```

7.6.2. Enabling the Kafka Bridge Grafana dashboard

If you deployed Kafka Bridge with your Kafka cluster, you can enable Grafana to present the metrics data it exposes.

A Kafka Bridge dashboard is provided in the **examples/metrics** directory as a JSON file:

- **strimzi-kafka-bridge.json**

When metrics data has been collected for some time, the Kafka Bridge charts are populated.

Kafka Bridge

Shows metrics for:

- HTTP connections to the Kafka Bridge count
- HTTP requests being processed count
- Requests processed per second grouped by HTTP method
- The total request rate grouped by response codes (2XX, 4XX, 5XX)
- Bytes received and sent per second
- Requests for each Kafka Bridge endpoint
- Number of Kafka consumers, producers, and related opened connections used by the Kafka Bridge itself
- Kafka producer:
  - The average number of records sent per second (grouped by topic)
  - The number of outgoing bytes sent to all brokers per second (grouped by topic)
  - The average number of records per second that resulted in errors (grouped by topic)
- Kafka consumer:
  - The average number of records consumed per second (grouped by clientId-topic)
  - The average number of bytes consumed per second (grouped by clientId-topic)
  - Partitions assigned (grouped by clientId)
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count

7.7. MONITOR CRUISE CONTROL

If you are already using Prometheus and Grafana for monitoring of built-in Kafka metrics, you can configure Prometheus to also scrape the Cruise Control Prometheus endpoint.

The example Grafana dashboard for Cruise Control provides:

- Information about optimization proposals computation, goals violation, cluster balancedness, and more
- Information about REST API calls for rebalance proposals and actual rebalance operations

- JVM metrics from Cruise Control itself

7.7.1. Configuring Cruise Control

Enable Cruise Control metrics using the **cruiseControl.metricsConfig** property in the **Kafka** resource to provide a reference to a ConfigMap that contains JMX exporter configuration for the metrics to expose.

For example:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
metadata:
  name: my-cluster
spec:
  # ...
  kafka:
    # ...
  zookeeper:
    # ...
  cruiseControl:
    metricsConfig:
      type: jmxPrometheusExporter
      valueFrom:
        configMapKeyRef:
          name: my-config-map
          key: my-key
```

7.7.2. Enabling the Cruise Control Grafana dashboard

If you deployed Cruise Control with your Kafka cluster with the metrics enabled, you can enable Grafana to present the metrics data it exposes.

A Cruise Control dashboard is provided in the **examples/metrics** directory as a JSON file:

- **strimzi-cruise-control.json**

When metrics data has been collected for some time, the Cruise Control charts are populated.

Cruise Control

Shows metrics for:

- Number of snapshot windows that are monitored by Cruise Control
- Number of time windows considered valid because they contain enough samples to compute an optimization proposal
- Number of ongoing executions running for proposals or rebalances
- Current balancedness score of the Kafka cluster as calculated by the anomaly detector component of Cruise Control (every 5 minutes by default)
- Percentage of monitored partitions
- Number of goal violations reported by the anomaly detector (every 5 minutes by default)

- How often a disk read failure happens on the brokers
- Rate of metric sample fetch failures
- Time needed to compute an optimization proposal
- Time needed to create the cluster model
- How often a proposal request or an actual rebalance request is made through the Cruise Control REST API
- How often the overall cluster state and the user tasks state are requested through the Cruise Control REST API
- JVM memory used
- JVM garbage collection time
- JVM garbage collection count

CHAPTER 8. UPGRADING AMQ STREAMS

AMQ Streams on OpenShift can be upgraded to version 1.7 to take advantage of new features and enhancements, performance improvements, and security options.

During this upgrade, you upgrade Kafka to the latest supported version. Each Kafka release introduces new features, improvements, and bug fixes to your AMQ Streams deployment.

AMQ Streams can be downgraded to the previous version if you encounter issues with the newer version.

Released versions of AMQ Streams are listed in the Product Downloads section of the Red Hat Customer Portal.

Upgrade paths

Two upgrade paths are possible:

Incremental

Upgrading AMQ Streams from the previous minor version to version 1.7.

Multi-version

Upgrading AMQ Streams from an old version to version 1.7 within a single upgrade (skipping one or more intermediate versions).
For example, upgrading from AMQ Streams 1.5 directly to 1.7.

Kafka version support

The Kafka versions table lists the supported Kafka versions for AMQ Streams 1.7. In the table:

- The *latest* Kafka version is supported for production use.
- The *previous* Kafka version is supported only for the purpose of upgrading to AMQ Streams 1.7.

Identify the Kafka version to upgrade to before you begin the upgrade procedures described in this chapter.

NOTE

You can upgrade to a higher Kafka version as long as it is supported by your version of AMQ Streams. In some cases, you can also downgrade to a previous supported Kafka version.

Downtime and availability

If topics are configured for high availability, upgrading AMQ Streams should not cause any downtime for consumers and producers that publish and read data from those topics. Highly available topics have a replication factor of at least 3 and partitions distributed evenly among the brokers.

Upgrading AMQ Streams triggers rolling updates, where all brokers are restarted in turn, at different stages of the process. During rolling updates, not all brokers are online, so overall *cluster availability* is temporarily reduced. A reduction in cluster availability increases the chance that a broker failure will result in lost messages.

8.1. AMQ STREAMS AND KAFKA UPGRADES

Upgrading AMQ Streams is a three-stage process. To upgrade brokers and clients without downtime, you *must* complete the upgrade procedures in the following order:

1. Update your Cluster Operator to a new AMQ Streams version.
   The approach you take depends on how you deployed the Cluster Operator.
   - If you deployed the Cluster Operator using the installation YAML files, perform your upgrade by modifying the Operator installation files, as described in Upgrading the Cluster Operator.
   - If you deployed the Cluster Operator from the OperatorHub, use the Operator Lifecycle Manager (OLM) to change the update channel for the AMQ Streams Operators to a new AMQ Streams version.
     Depending on your chosen upgrade strategy, after updating the channel, either:
     - An automatic upgrade is initiated
     - A manual upgrade will require approval before the installation begins
       For more information on using the OperatorHub to upgrade Operators, see Upgrading installed Operators in the OpenShift documentation.
2. Upgrade all Kafka brokers and client applications to the latest supported Kafka version.
   - Section 8.1.3, “Upgrading Kafka”
   - Section 8.1.5, “Strategies for upgrading clients”
3. If applicable, perform the following tasks:
   a. Update existing custom resources to handle deprecated custom resource properties.
      - Section 8.2, “AMQ Streams custom resource upgrades”

      

      NOTE

      Custom resources can also be updated *before* the Kafka upgrade.

   b. Update listeners to use the **GenericKafkaListener** schema
      - Section 8.1.4, “Updating listeners to the generic listener configuration”

Optional: incremental cooperative rebalance upgrade

Consider upgrading consumers and Kafka Streams applications to use the *incremental cooperative rebalance* protocol for partition rebalances.

- Section 8.3, “Upgrading consumers to cooperative rebalancing”

8.1.1. Kafka versions

Kafka’s log message format version and inter-broker protocol version specify, respectively, the log format version appended to messages and the version of the Kafka protocol used in a cluster. To ensure the correct versions are used, the upgrade process involves making configuration changes to existing Kafka brokers and code changes to client applications (consumers and producers).

The following table shows the differences between Kafka versions:

Kafka version	Interbroker protocol version	Log message format version	ZooKeeper version
2.6.0	2.6	2.6	3.5.8
2.7.0	2.7	2.7	3.5.8

Inter-broker protocol version

In Kafka, the network protocol used for inter-broker communication is called the *inter-broker protocol*. Each version of Kafka has a compatible version of the inter-broker protocol. The minor version of the protocol typically increases to match the minor version of Kafka, as shown in the preceding table.

The inter-broker protocol version is set cluster wide in the **Kafka** resource. To change it, you edit the **inter.broker.protocol.version** property in **Kafka.spec.kafka.config**.

Log message format version

When a producer sends a message to a Kafka broker, the message is encoded using a specific format. The format can change between Kafka releases, so messages specify which version of the format they were encoded with. You can configure a Kafka broker to convert messages from newer format versions to a given older format version before the broker appends the message to the log.

In Kafka, there are two different methods for setting the message format version:

- The **message.format.version** property is set on topics.
- The **log.message.format.version** property is set on Kafka brokers.

The default value of **message.format.version** for a topic is defined by the **log.message.format.version** that is set on the Kafka broker. You can manually set the **message.format.version** of a topic by modifying its topic configuration.

The upgrade tasks in this section assume that the message format version is defined by the **log.message.format.version**.

8.1.2. Upgrading the Cluster Operator

The steps to upgrade your Cluster Operator deployment to use AMQ Streams 1.7 are described in this section.

Follow this procedure if you deployed the Cluster Operator using the installation YAML files rather than OperatorHub.

The availability of Kafka clusters managed by the Cluster Operator is not affected by the upgrade operation.

NOTE

Refer to the documentation supporting a specific version of AMQ Streams for information on how to upgrade to that version.

8.1.2.1. Upgrading the Cluster Operator

This procedure describes how to upgrade a Cluster Operator deployment to use AMQ Streams 1.7.

Prerequisites

- An existing Cluster Operator deployment is available.
- You have downloaded the release artifacts for AMQ Streams 1.7 .

Procedure

1. Take note of any configuration changes made to the existing Cluster Operator resources (in the **/install/cluster-operator** directory). Any changes will be **overwritten** by the new version of the Cluster Operator.
2. Update your custom resources to reflect the supported configuration options available for AMQ Streams version 1.7.
3. Update the Cluster Operator.
   a. Modify the installation files for the new Cluster Operator version according to the namespace the Cluster Operator is running in.
   On Linux, use:

   ```
   sed -i 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

   On MacOS, use:

   ```
   sed -i '' 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

   b. If you modified one or more environment variables in your existing Cluster Operator **Deployment**, edit the **install/cluster-operator/060-Deployment-strimzi-cluster-operator.yaml** file to use those environment variables.
4. When you have an updated configuration, deploy it along with the rest of the installation resources:

   ```
   oc replace -f install/cluster-operator
   ```

   Wait for the rolling updates to complete.
5. If the new Operator version no longer supports the Kafka version you are upgrading from, the Cluster Operator returns a "Version not found" error message. Otherwise, no error message is returned.
   For example:

   ```
   "Version 2.4.0 is not supported. Supported versions are: 2.6.0, 2.6.1, 2.7.0."
   ```

   - If the error message is returned, upgrade to a Kafka version that is supported by the new Cluster Operator version:
     a. Edit the **Kafka** custom resource.
     b. Change the **spec.kafka.version** property to a supported Kafka version.

- If the error message is *not* returned, go to the next step. You will upgrade the Kafka version later.

6. Get the image for the Kafka pod to ensure the upgrade was successful:

```
oc get pods my-cluster-kafka-0 -o jsonpath='{.spec.containers[0].image}'
```

The image tag shows the new Operator version. For example:

```
registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:{ContainerVersion}
```

Your Cluster Operator was upgraded to version 1.7 but the version of Kafka running in the cluster it manages is unchanged.

Following the Cluster Operator upgrade, you must perform a Kafka upgrade.

8.1.3. Upgrading Kafka

After you have upgraded your Cluster Operator to 1.7, the next step is to upgrade all Kafka brokers to the latest supported version of Kafka.

Kafka upgrades are performed by the Cluster Operator through rolling updates of the Kafka brokers.

The Cluster Operator initiates rolling updates based on the Kafka cluster configuration.

If **Kafka.spec.kafka.config** contains...	The Cluster Operator initiates...
Both the **inter.broker.protocol.version** and the **log.message.format.version**.	A single rolling update. After the update, the **inter.broker.protocol.version** must be updated manually, followed by **log.message.format.version**. Changing each will trigger a further rolling update.
Either the **inter.broker.protocol.version** or the **log.message.format.version**.	Two rolling updates.
No configuration for the **inter.broker.protocol.version** or the **log.message.format.version**.	Two rolling updates.

As part of the Kafka upgrade, the Cluster Operator initiates rolling updates for ZooKeeper.

- A single rolling update occurs even if the ZooKeeper version is unchanged.
- Additional rolling updates occur if the new version of Kafka requires a new ZooKeeper version.

Additional resources

- Section 8.1.2, “Upgrading the Cluster Operator”
- Section 8.1.1, “Kafka versions”

8.1.3.1. Kafka version and image mappings

When upgrading Kafka, consider your settings for the **STRIMZI_KAFKA_IMAGES** environment variable and the **Kafka.spec.kafka.version** property.

- Each **Kafka** resource can be configured with a **Kafka.spec.kafka.version**.
- The Cluster Operator's **STRIMZI_KAFKA_IMAGES** environment variable provides a mapping between the Kafka version and the image to be used when that version is requested in a given **Kafka** resource.
  - If **Kafka.spec.kafka.image** is not configured, the default image for the given version is used.
  - If **Kafka.spec.kafka.image** is configured, the default image is overridden.

WARNING

The Cluster Operator cannot validate that an image actually contains a Kafka broker of the expected version. Take care to ensure that the given image corresponds to the given Kafka version.

8.1.3.2. Upgrading Kafka brokers and client applications

This procedure describes how to upgrade a AMQ Streams Kafka cluster to the latest supported Kafka version.

Compared to your current Kafka version, the new version might support a higher *log message format version* or *inter-broker protocol version*, or both. Follow the steps to upgrade these versions, if required. For more information, see Section 8.1.1, "Kafka versions".

You should also choose a strategy for upgrading clients. Kafka clients are upgraded in step 6 of this procedure.

Prerequisites

For the **Kafka** resource to be upgraded, check that:

- The Cluster Operator, which supports both versions of Kafka, is up and running.
- The **Kafka.spec.kafka.config** does *not* contain options that are not supported in the new Kafka version.

Procedure

1. Update the Kafka cluster configuration:

```
oc edit kafka my-cluster
```

2. If configured, ensure that **Kafka.spec.kafka.config** has the **log.message.format.version** and **inter.broker.protocol.version** set to the defaults for the *current* Kafka version.
   For example, if upgrading from Kafka version 2.6.0 to 2.7.0:

```
kind: Kafka
spec:
  # ...
  kafka:
    version: 2.6.0
    config:
      log.message.format.version: "2.6"
      inter.broker.protocol.version: "2.6"
      # ...
```

If **log.message.format.version** and **inter.broker.protocol.version** are not configured, AMQ Streams automatically updates these versions to the current defaults after the update to the Kafka version in the next step.

NOTE

The value of **log.message.format.version** and **inter.broker.protocol.version** must be strings to prevent them from being interpreted as floating point numbers.

3. Change the **Kafka.spec.kafka.version** to specify the new Kafka version; leave the **log.message.format.version** and **inter.broker.protocol.version** at the defaults for the *current* Kafka version.

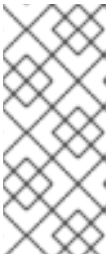

NOTE

Changing the **kafka.version** ensures that all brokers in the cluster will be upgraded to start using the new broker binaries. During this process, some brokers are using the old binaries while others have already upgraded to the new ones. Leaving the **inter.broker.protocol.version** unchanged ensures that the brokers can continue to communicate with each other throughout the upgrade.

For example, if upgrading from Kafka 2.6.0 to 2.7.0:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
spec:
  # ...
  kafka:
    version: 2.7.0 (1)
    config:
      log.message.format.version: "2.6" (2)
      inter.broker.protocol.version: "2.6" (3)
      # ...
```

(1) Kafka version is changed to the new version.

(2) Message format version is unchanged.

(3) Inter-broker protocol version is unchanged.

WARNING

You cannot downgrade Kafka if the **inter.broker.protocol.version** for the new Kafka version changes. The inter-broker protocol version determines the schemas used for persistent metadata stored by the broker, including messages written to **__consumer_offsets**. The downgraded cluster will not understand the messages.

4. If the image for the Kafka cluster is defined in the Kafka custom resource, in **Kafka.spec.kafka.image**, update the **image** to point to a container image with the new Kafka version.
   See Kafka version and image mappings

5. Save and exit the editor, then wait for rolling updates to complete.
   Check the progress of the rolling updates by watching the pod state transitions:

   ```
   oc get pods my-cluster-kafka-0 -o jsonpath='{.spec.containers[0].image}'
   ```

   The rolling updates ensure that each pod is using the broker binaries for the new version of Kafka.

6. Depending on your chosen strategy for upgrading clients, upgrade all client applications to use the new version of the client binaries.
   If required, set the **version** property for Kafka Connect and MirrorMaker as the new version of Kafka:

   a. For Kafka Connect, update **KafkaConnect.spec.version**.

   b. For MirrorMaker, update **KafkaMirrorMaker.spec.version**.

   c. For MirrorMaker 2.0, update **KafkaMirrorMaker2.spec.version**.

7. If configured, update the Kafka resource to use the new **inter.broker.protocol.version** version.
   Otherwise, go to step 9.
   For example, if upgrading to Kafka 2.7.0:

   ```
   apiVersion: kafka.strimzi.io/v1beta2
   kind: Kafka
   spec:
     # ...
     kafka:
       version: 2.7.0
       config:
         log.message.format.version: "2.6"
         inter.broker.protocol.version: "2.7"
         # ...
   ```

8. Wait for the Cluster Operator to update the cluster.

9. If configured, update the Kafka resource to use the new **log.message.format.version** version.
   Otherwise, go to step 10.

For example, if upgrading to Kafka 2.7.0:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
spec:
  # ...
  kafka:
    version: 2.7.0
    config:
      log.message.format.version: "2.7"
      inter.broker.protocol.version: "2.7"
      # ...
```

10. Wait for the Cluster Operator to update the cluster.

    - The Kafka cluster and clients are now using the new Kafka version.
    - The brokers are configured to send messages using the inter-broker protocol version and message format version of the new version of Kafka.

Following the Kafka upgrade, if required, you can:

- Update listeners to the **GenericKafkaListener** schema
- Upgrade consumers to use the incremental cooperative rebalance protocol
- Update existing custom resources

8.1.4. Updating listeners to the generic listener configuration

AMQ Streams provides a **GenericKafkaListener** schema for the configuration of Kafka listeners in a **Kafka** resource.

GenericKafkaListener replaces the **KafkaListeners** schema, which has been removed from AMQ Streams.

With the **GenericKafkaListener** schema, you can configure as many listeners as required, as long as their names and ports are unique. The **listeners** configuration is defined as an array, but the deprecated format is also supported.

For clients inside the OpenShift cluster, you can create **plain** (without encryption) or **tls** *internal* listeners.

For clients outside the OpenShift cluster, you create *external* listeners and specify a connection mechanism, which can be **nodeport**, **loadbalancer**, **ingress** or **route**.

The **KafkaListeners** schema used sub-properties for **plain**, **tls** and **external** listeners, with fixed ports for each. At any stage in the upgrade process, you must convert listeners configured using the **KafkaListeners** schema into the format of the **GenericKafkaListener** schema.

For example, if you are currently using the following configuration in your **Kafka** configuration:

Old listener configuration

```
listeners:
  plain:
```

```
    # ...
  tls:
    # ...
  external:
    type: loadbalancer
    # ...
```

Convert the listeners into the new format using:

New listener configuration

```
listeners:
  #...
  - name: plain
    port: 9092
    type: internal
    tls: false (1)
  - name: tls
    port: 9093
    type: internal
    tls: true
  - name: external
    port: 9094
    type: EXTERNAL-LISTENER-TYPE (2)
    tls: true
```

1. The TLS property is now required for all listeners.
2. Options: **ingress**, **loadbalancer**, **nodeport**, **route**.

Make sure to use the **exact** names and port numbers shown.

For any additional **configuration** or **overrides** properties used with the old format, you need to update them to the new format.

Changes introduced to the listener **configuration**:

- **overrides** is merged with the **configuration** section
- **dnsAnnotations** has been renamed **annotations**
- **preferredAddressType** has been renamed **preferredNodePortAddressType**
- **address** has been renamed **alternativeNames**
- **loadBalancerSourceRanges** and **externalTrafficPolicy** move to the listener configuration from the now deprecated **template**

For example, this configuration:

Old additional listener configuration

```
listeners:
  external:
    type: loadbalancer
```

```
    authentication:
      type: tls
    overrides:
      bootstrap:
        dnsAnnotations:
          #...
```

Changes to:

New additional listener configuration

```
listeners:
    #...
  - name: external
    port: 9094
    type:loadbalancer
    tls: true
    authentication:
      type: tls
    configuration:
      bootstrap:
        annotations:
          #...
```

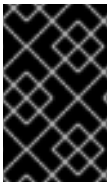

IMPORTANT

The name and port numbers shown in the new listener configuration **must** be used for backwards compatibility. Using any other values will cause renaming of the Kafka listeners and OpenShift services.

For more information on the configuration options available for each type of listener, see the **GenericKafkaListener** schema reference.

8.1.5. Strategies for upgrading clients

The right approach to upgrading your client applications (including Kafka Connect connectors) depends on your particular circumstances.

Consuming applications need to receive messages in a message format that they understand. You can ensure that this is the case in one of two ways:

- By upgrading all the consumers for a topic *before* upgrading any of the producers.
- By having the brokers down-convert messages to an older format.

Using broker down-conversion puts extra load on the brokers, so it is not ideal to rely on down-conversion for all topics for a prolonged period of time. For brokers to perform optimally they should not be down converting messages at all.

Broker down-conversion is configured in two ways:

- The topic-level **message.format.version** configures it for a single topic.
- The broker-level **log.message.format.version** is the default for topics that do not have the topic-level **message.format.version** configured.

Messages published to a topic in a new-version format will be visible to consumers, because brokers perform down-conversion when they receive messages from producers, not when they are sent to consumers.

There are a number of strategies you can use to upgrade your clients:

Consumers first

1. Upgrade all the consuming applications.
2. Change the broker-level **log.message.format.version** to the new version.
3. Upgrade all the producing applications.
   This strategy is straightforward, and avoids any broker down-conversion. However, it assumes that all consumers in your organization can be upgraded in a coordinated way, and it does not work for applications that are both consumers and producers. There is also a risk that, if there is a problem with the upgraded clients, new-format messages might get added to the message log so that you cannot revert to the previous consumer version.

Per-topic consumers first

For each topic:

1. Upgrade all the consuming applications.
2. Change the topic-level **message.format.version** to the new version.
3. Upgrade all the producing applications.
   This strategy avoids any broker down-conversion, and means you can proceed on a topic-by-topic basis. It does not work for applications that are both consumers and producers of the same topic. Again, it has the risk that, if there is a problem with the upgraded clients, new-format messages might get added to the message log.

Per-topic consumers first, with down conversion

For each topic:

1. Change the topic-level **message.format.version** to the old version (or rely on the topic defaulting to the broker-level **log.message.format.version**).
2. Upgrade all the consuming and producing applications.
3. Verify that the upgraded applications function correctly.
4. Change the topic-level **message.format.version** to the new version.
   This strategy requires broker down-conversion, but the load on the brokers is minimized because it is only required for a single topic (or small group of topics) at a time. It also works for applications that are both consumers and producers of the same topic. This approach ensures that the upgraded producers and consumers are working correctly before you commit to using the new message format version.

   The main drawback of this approach is that it can be complicated to manage in a cluster with many topics and applications.

Other strategies for upgrading client applications are also possible.

NOTE

It is also possible to apply multiple strategies. For example, for the first few applications and topics the "per-topic consumers first, with down conversion" strategy can be used. When this has proved successful another, more efficient strategy can be considered acceptable to use instead.

8.2. AMQ STREAMS CUSTOM RESOURCE UPGRADES

After you have upgraded AMQ Streams to 1.7, you must ensure that your custom resources are using API version **v1beta2**. You can do this any time after upgrading to 1.7, but the upgrades **must be completed before the next AMQ Streams minor version update.**

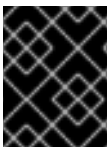

IMPORTANT

Upgrade of the custom resources to **v1beta2 must** be performed after upgrading the Cluster Operator, so the Cluster Operator can understand the resources.

NOTE

Upgrade of the custom resources to **v1beta2** prepares AMQ Streams for a move to OpenShift CRD **v1**, which is required for Kubernetes v1.22.

CLI upgrades to custom resources

AMQ Streams provides an *API conversion tool* with its release artifacts.

You can download its ZIP or TAR.GZ from AMQ Streams download site . To use the tool, extract it and use the scripts in the **bin** directory.

From its CLI, you can then use the tool to convert the format of your custom resources to **v1beta2** in one of two ways:

- Section 8.2.2, “Converting custom resources configuration files using the API conversion tool”
- Section 8.2.3, “Converting custom resources directly using the API conversion tool”

After the conversion of your custom resources, you must set **v1beta2** as the *storage* API version in your CRDs:

- Section 8.2.4, “Upgrading CRDs to v1beta2 using the API conversion tool”

Manual upgrades to custom resources

Instead of using the API conversion tool to update custom resources to **v1beta2**, you can manually update each custom resource to use **v1beta2**:

Update the **Kafka** custom resource, including the configurations for the other components:

- Section 8.2.5, “Upgrading Kafka resources to support v1beta2”
- Section 8.2.6, “Upgrading ZooKeeper to support v1beta2”
- Section 8.2.7, “Upgrading the Topic Operator to support v1beta2”
- Section 8.2.8, “Upgrading the Entity Operator to support v1beta2”

- Section 8.2.9, “Upgrading Cruise Control to support v1beta2” (if Cruise Control is deployed)
- Section 8.2.10, “Upgrading the API version of Kafka resources to v1beta2”

Update the other custom resources that apply to your deployment:

- Section 8.2.11, “Upgrading Kafka Connect resources to v1beta2”
- Section 8.2.12, “Upgrading Kafka Connect S2I resources to v1beta2”
- Section 8.2.13, “Upgrading Kafka MirrorMaker resources to v1beta2”
- Section 8.2.14, “Upgrading Kafka MirrorMaker 2.0 resources to v1beta2”
- Section 8.2.15, “Upgrading Kafka Bridge resources to v1beta2”
- Section 8.2.16, “Upgrading Kafka User resources to v1beta2”
- Section 8.2.17, “Upgrading Kafka Topic resources to v1beta2”
- Section 8.2.18, “Upgrading Kafka Connector resources to v1beta2”
- Section 8.2.19, “Upgrading Kafka Rebalance resources to v1beta2”

The manual procedures show the changes that are made to each custom resource. After these changes, you must use the API conversion tool to upgrade your CRDs.

8.2.1. API versioning

Custom resources are edited and controlled using APIs added to OpenShift by CRDs. Put another way, CRDs extend the Kubernetes API to allow the creation of custom resources. CRDs are themselves resources within OpenShift. They are installed in an OpenShift cluster to define the versions of API for the custom resource. Each version of the custom resource API can define its own schema for that version. OpenShift clients, including the AMQ Streams Operators, access the custom resources served by the Kubernetes API server using a URL path (*API path*), which includes the API version.

The introduction of **v1beta2** updates the schemas of the custom resources. Older API versions are deprecated.

The **v1alpha1** API version is deprecated for the following AMQ Streams custom resources:

- **Kafka**
- **KafkaConnect**
- **KafkaConnectS2I**
- **KafkaConnector**
- **KafkaMirrorMaker**
- **KafkaMirrorMaker2**
- **KafkaTopic**
- **KafkaUser**

- **KafkaBridge**
- **KafkaRebalance**

The **v1beta1** API version is deprecated for the following AMQ Streams custom resources:

- **Kafka**
- **KafkaConnect**
- **KafkaConnectS2I**
- **KafkaMirrorMaker**
- **KafkaTopic**
- **KafkaUser**

IMPORTANT

The **v1alpha1** and **v1beta1** versions will be removed in the next minor release.

Additional resources

- Extend the Kubernetes API with CustomResourceDefinitions

8.2.2. Converting custom resources configuration files using the API conversion tool

This procedure describes how to use the API conversion tool to convert YAML files describing the configuration for AMQ Streams custom resources into a format applicable to **v1beta2**. To do so, you use the **convert-file** (**cf**) command.

The **convert-file** command can convert YAML files containing multiple documents. For a multi-document YAML file, all the AMQ Streams custom resources it contains are converted. Any non-AMQ Streams OpenShift resources are replicated unmodified in the converted output file.

After you have converted the YAML file, you must apply the configuration to update the custom resource in the cluster. Alternatively, if the GitOps synchronization mechanism is being used for updates on your cluster, you can use it to apply the changes. The conversion is only complete when the custom resource is updated in the OpenShift cluster.

Alternatively, you can use the **convert-resource** procedure to convert custom resources directly .

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- The API conversion tool, which is provided with the release artifacts.
- The tool requires Java 11.

Use the CLI help for more information on the API conversion tool, and the flags available for the **convert-file** command:

```
bin/api-conversion.sh help
bin/api-conversion.sh help convert-file
```

Use **bin/api-conversion.cmd** for this procedure if you are using Windows.

Table 8.1. Flags for YAML file conversion

Flag	Description
-f, **--file=*NAME-OF-YAML-FILE***	Specifies the YAML file for the AMQ Streams custom resource being converted
-o, --output=*NAME-OF-CONVERTED-YAML-FILE*	Creates an output YAML file for the converted custom resource
--in-place	Updates the original source file with the converted YAML

Procedure

1. Run the API conversion tool with the **convert-file** command and appropriate flags.
   Example 1, converts a YAML file and displays the output, though the file does not change:

   ```
   bin/api-conversion.sh convert-file --file input.yaml
   ```

   Example 2, converts a YAML file, and writes the changes into the original source file:

   ```
   bin/api-conversion.sh convert-file --file input.yaml --in-place
   ```

   Example 3, converts a YAML file, and writes the changes into a new output file:

   ```
   bin/api-conversion.sh convert-file --file input.yaml --output output.yaml
   ```

2. Update the custom resources using the converted configuration file.

   ```
   oc apply -f CONVERTED-CONFIG-FILE
   ```

3. Verify that the custom resources have been converted.

   ```
   oc get KIND CUSTOM-RESOURCE-NAME -o yaml
   ```

8.2.3. Converting custom resources directly using the API conversion tool

This procedure describes how to use the API conversion tool to convert AMQ Streams custom resources directly in the OpenShift cluster into a format applicable to **v1beta2**. To do so, you use the **convert-resource** (**cr**) command. The command uses Kubernetes APIs to make the conversions.

You can specify one or more of types of AMQ Streams custom resources, based on the **kind** property, or you can convert all types. You can also target a specific namespace or all namespaces for conversion. When targeting a namespace, you can convert all custom resources in that namespace, or convert a single custom resource by specifying its name and kind.

Alternatively, you can use the **convert-file** procedure to convert and apply the YAML files describing the custom resources.

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- The API conversion tool, which is provided with the release artifacts.
- The tool requires Java 11 (OpenJDK).
- The steps require a user admin account with RBAC permission to:
  - Get the AMQ Streams custom resources being converted using the **--name** option
  - List the AMQ Streams custom resources being converted without using the **--name** option
  - Replace the AMQ Streams custom resources being converted

Use the CLI help for more information on the API conversion tool, and the flags available for the **convert-resource** command:

```
bin/api-conversion.sh help
bin/api-conversion.sh help convert-resource
```

Use **bin/api-conversion.cmd** for this procedure if you are using Windows.

Table 8.2. Flags for converting custom resources

Flag	Description
-k, **--kind**	Specifies the kinds of custom resources to be converted, or converts all resources if not specified
-a, **--all-namespaces**	Converts custom resources in all namespaces
-n, **--namespace**	Specifies an OpenShift namespace or OpenShift project, or uses the current namespace if not specified
--name	If **--namespace** and a single custom resource**--kind** is used, specifies the name of the custom resource being converted

Procedure

1. Run the API conversion tool with the **convert-resource** command and appropriate flags.
   Example 1, converts all AMQ Streams resources in current namespace:

   ```
   bin/api-conversion.sh convert-resource
   ```

   Example 2, converts all AMQ Streams resources in all namespaces:

```
bin/api-conversion.sh convert-resource --all-namespaces
```

Example 3, converts all AMQ Streams resources in the **my-kafka** namespace:

```
bin/api-conversion.sh convert-resource --namespace my-kafka
```

Example 4, converts only Kafka resources in all namespaces:

```
bin/api-conversion.sh convert-resource --all-namespaces --kind Kafka
```

Example 5, converts Kafka and Kafka Connect resources in all namespaces:

```
bin/api-conversion.sh convert-resource --all-namespaces --kind Kafka --kind KafkaConnect
```

Example 6, converts a Kafka custom resource named **my-cluster** in the **my-kafka** namespace:

```
bin/api-conversion.sh convert-resource --kind Kafka --namespace my-kafka --name my-cluster
```

2. Verify that the custom resources have been converted.

```
oc get KIND CUSTOM-RESOURCE-NAME -o yaml
```

8.2.4. Upgrading CRDs to v1beta2 using the API conversion tool

This procedure describes how to use the API conversion tool to convert the CRDs that define the schemas used to instantiate and manage AMQ Streams-specific resources in a format applicable to **v1beta2**. To do so, you use the **crd-upgrade** command.

Perform this procedure after converting all AMQ Streams custom resources in the whole OpenShift cluster to v1beta2. If you upgrade your CRDs first, and then convert your custom resources, you will need to run this command again.

The command updates **spec.versions** in the CRDs to declare **v1beta2** as the *storage* API version. The command also updates custom resources so they are stored under **v1beta2**. New custom resource instances are created from the specification of the storage API version, so only one API version is ever marked as the storage version.

When you have upgraded the CRDs to use **v1beta2** as the storage version, you should only use **v1beta2** properties in your custom resources.

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- The API conversion tool, which is provided with the release artifacts.
- The tool requires Java 11 (OpenJDK).
- Custom resources have been converted to **v1beta2**.
- The steps require a user admin account with RBAC permission to:

- List the AMQ Streams custom resources in all namespaces
- Replace the AMQ Streams custom resources being converted
- Update CRDs
- Replace the status of the CRDs

Use the CLI help for more information on the API conversion tool:

```
bin/api-conversion.sh help
```

Use **bin/api-conversion.cmd** for this procedure if you are using Windows.

Procedure

1. If you have not done so, convert your custom resources to use **v1beta2**.
   You can use the API conversion tool to do this in one of two ways:
   - Section 8.2.2, "Converting custom resources configuration files using the API conversion tool"
   - Section 8.2.3, "Converting custom resources directly using the API conversion tool"
     Or you can make the changes manually.
2. Run the API conversion tool with the **crd-upgrade** command.

   ```
   bin/api-conversion.sh crd-upgrade
   ```
3. Verify that the CRDs have been upgraded so that v1beta2 is the storage version.
   For example, for the Kafka topic CRD:

   ```
   apiVersion: kafka.strimzi.io/v1beta2
   kind: CustomResourceDefinition
   metadata:
     name: kafkatopics.kafka.strimzi.io
     #...
   spec:
     group: kafka.strimzi.io
     #...
     versions:
     - name: v1beta2
       served: true
       storage: true
       #...
   status:
     #...
     storedVersions:
     - v1beta2
   ```

8.2.5. Upgrading Kafka resources to support v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **Kafka** custom resource in your deployment.

1. Update the **Kafka** custom resource in an editor.

```
oc edit kafka KAFKA-CLUSTER
```

2. If you have not already done so, update **.spec.kafka.listener** to the new generic listener format, as described in Section 8.1.4, "Updating listeners to the generic listener configuration" .

WARNING

The old listener format is not supported in API version **v1beta2**.

3. If present, move **affinity** from **.spec.kafka.affinity** to **.spec.kafka.template.pod.affinity**.
4. If present, move **tolerations** from **.spec.kafka.tolerations** to **.spec.kafka.template.pod.tolerations**.
5. If present, remove **.spec.kafka.template.tlsSidecarContainer**.
6. If present, remove **.spec.kafka.tlsSidecarContainer**.
7. If either of the following policy configurations exist:
   - **.spec.kafka.template.externalBootstrapService.externalTrafficPolicy**
   - **.spec.kafka.template.perPodService.externalTrafficPolicy**
     a. Move the configuration to **.spec.kafka.listeners[].configuration.externalTrafficPolicy**, for both **type: loadbalancer** and **type: nodeport** listeners.
     b. Remove **.spec.kafka.template.externalBootstrapService.externalTrafficPolicy** or **.spec.kafka.template.perPodService.externalTrafficPolicy**.
8. If either of the following **loadbalancer** listener configurations exist:
   - **.spec.kafka.template.externalBootstrapService.loadBalancerSourceRanges**
   - **.spec.kafka.template.perPodService.loadBalancerSourceRanges**
     a. Move the configuration to **.spec.kafka.listeners[].configuration.loadBalancerSourceRanges**, for **type: loadbalancer** listeners.
     b. Remove **.spec.kafka.template.externalBootstrapService.loadBalancerSourceRanges** or **.spec.kafka.template.perPodService.loadBalancerSourceRanges**.

9. If **type: external** logging is configured in **.spec.kafka.logging**:
   Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

   With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j.properties
```

10. If the **.spec.kafka.metrics** field is used to enable metrics:

    a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.kafka.metrics** field.

```
kind: ConfigMap
apiVersion: v1
metadata:
  name: kafka-metrics
  labels:
    app: strimzi
data:
  kafka-metrics-config.yaml: |
    <YAML>
```

    b. Add a **.spec.kafka.metricsConfig** property that points to the ConfigMap and key:

```
metricsConfig:
  type: jmxPrometheusExporter
  valueFrom:
    configMapKeyRef:
      name: kafka-metrics
      key: kafka-metrics-config.yaml
```

    c. Delete the old **.spec.kafka.metrics** field.

11. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

What to do next

For each **Kafka** custom resource, upgrade the configurations for ZooKeeper, Topic Operator, Entity Operator, and Cruise Control (if deployed) to support version **v1beta2**. This is described in the following procedures.

When all **Kafka** configurations are updated to support **v1beta2**, you can upgrade the **Kafka** custom resource to **v1beta2**.

8.2.6. Upgrading ZooKeeper to support v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **Kafka** custom resource in your deployment.

1. Update the **Kafka** custom resource in an editor.

```
oc edit kafka KAFKA-CLUSTER
```

2. If present, move **affinity** from **.spec.zookeeper.affinity** to **.spec.zookeeper.template.pod.affinity**.

3. If present, move **tolerations** from **.spec.zookeeper.tolerations** to **.spec.zookeeper.template.pod.tolerations**.

4. If present, remove **.spec.zookeeper.template.tlsSidecarContainer**.

5. If present, remove **.spec.zookeeper.tlsSidecarContainer**.

6. If **type: external** logging is configured in **.spec.kafka.logging**:
   Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

   With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j.properties
```

7. If the **.spec.zookeeper.metrics** field is used to enable metrics:

   a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.zookeeper.metrics** field.

```
kind: ConfigMap
apiVersion: v1
metadata:
  name: kafka-metrics
  labels:
    app: strimzi
```

```
data:
  zookeeper-metrics-config.yaml: |
    <YAML>
```

b. Add a **.spec.zookeeper.metricsConfig** property that points to the ConfigMap and key:

```
metricsConfig:
  type: jmxPrometheusExporter
  valueFrom:
    configMapKeyRef:
      name: kafka-metrics
      key: zookeeper-metrics-config.yaml
```

c. Delete the old **.spec.zookeeper.metrics** field.

8. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.7. Upgrading the Topic Operator to support v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **Kafka** custom resource in your deployment.

1. Update the **Kafka** custom resource in an editor.

```
oc edit kafka KAFKA-CLUSTER
```

2. If **Kafka.spec.topicOperator** is used:

   a. Move **affinity** from **.spec.topicOperator.affinity** to **.spec.entityOperator.template.pod.affinity**.

   b. Move **tolerations** from **.spec.topicOperator.tolerations** to **.spec.entityOperator.template.pod.tolerations**.

   c. Move **.spec.topicOperator.tlsSidecar** to **.spec.entityOperator.tlsSidecar**.

   d. After moving **affinity**, **tolerations**, and **tlsSidecar**, move the remaining configuration in **.spec.topicOperator** to **.spec.entityOperator.topicOperator**.

3. If **type: external** logging is configured in **.spec.topicOperator.logging**:
   Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

   With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
```

```
type: external
valueFrom:
  configMapKeyRef:
    name: my-config-map
    key: log4j2.properties
```

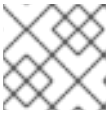

NOTE

You can also complete this step as part of the Entity Operator upgrade.

4. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.8. Upgrading the Entity Operator to support v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- **Kafka.spec.entityOperator** is configured, as described in Section 8.2.7, "Upgrading the Topic Operator to support v1beta2".

Procedure

Perform the following steps for each **Kafka** custom resource in your deployment.

1. Update the **Kafka** custom resource in an editor.

```
oc edit kafka KAFKA-CLUSTER
```

2. Move **affinity** from **.spec.entityOperator.affinity** to **.spec.entityOperator.template.pod.affinity**.

3. Move **tolerations** from **.spec.entityOperator.tolerations** to **.spec.entityOperator.template.pod.tolerations**.

4. If **type: external** logging is configured in **.spec.entityOperator.userOperator.logging** or **.spec.entityOperator.topicOperator.logging**:
Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j2.properties
```

5. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.9. Upgrading Cruise Control to support v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- Cruise Control is configured and deployed. See Deploying Cruise Control in the *Using AMQ Streams on OpenShift* guide.

Procedure

Perform the following steps for each **Kafka.spec.cruiseControl** configuration in your Kafka cluster.

1. Update the **Kafka** custom resource in an editor.

```
oc edit kafka KAFKA-CLUSTER
```

2. If **type: external** logging is configured in **.spec.cruiseControl.logging**:
Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j2.properties
```

3. If the **.spec.cruiseControl.metrics** field is used to enable metrics:

   a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.cruiseControl.metrics** field.

```
kind: ConfigMap
apiVersion: v1
metadata:
  name: kafka-metrics
  labels:
    app: strimzi
data:
  cruise-control-metrics-config.yaml: |
    <YAML>
```

   b. Add a **.spec.cruiseControl.metricsConfig** property that points to the ConfigMap and key:

```
metricsConfig:
  type: jmxPrometheusExporter
  valueFrom:
    configMapKeyRef:
      name: kafka-metrics
      key: cruise-control-metrics-config.yaml
```

c. Delete the old **.spec.cruiseControl.metrics** field.

4. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.10. Upgrading the API version of Kafka resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- You have updated the following configurations within the **Kafka** custom resource:
  - ZooKeeper
  - Topic Operator
  - Entity Operator
  - Cruise Control (if Cruise Control is deployed)

Procedure

Perform the following steps for each **Kafka** custom resource in your deployment.

1. Update the **Kafka** custom resource in an editor.

```
oc edit kafka KAFKA-CLUSTER
```

2. Update the **apiVersion** of the **Kafka** custom resource to **v1beta2**:
Replace:

```
apiVersion: kafka.strimzi.io/v1beta1
```

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

3. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.11. Upgrading Kafka Connect resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **KafkaConnect** custom resource in your deployment.

1. Update the **KafkaConnect** custom resource in an editor.

```
oc edit kafkaconnect KAFKA-CONNECT-CLUSTER
```

2. If present, move:

```
KafkaConnect.spec.affinity
```

```
KafkaConnect.spec.tolerations
```

to:

```
KafkaConnect.spec.template.pod.affinity
```

```
KafkaConnect.spec.template.pod.tolerations
```

For example, move:

```
spec:
  # ...
  affinity:
    # ...
  tolerations:
    # ...
```

to:

```
spec:
  # ...
  template:
    pod:
      affinity:
        # ...
      tolerations:
        # ...
```

3. If **type: external** logging is configured in **.spec.logging**:
Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
```

```
    configMapKeyRef:
      name: my-config-map
      key: log4j.properties
```

4. If the **.spec.metrics** field is used to enable metrics:

   a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.metrics** field.

```
kind: ConfigMap
apiVersion: v1
metadata:
  name: kafka-connect-metrics
  labels:
    app: strimzi
data:
  connect-metrics-config.yaml: |
    <YAML>
```

   b. Add a **.spec.metricsConfig** property that points to the ConfigMap and key:

```
metricsConfig:
  type: jmxPrometheusExporter
  valueFrom:
    configMapKeyRef:
      name: kafka-connect-metrics
      key: connect-metrics-config.yaml
```

   c. Delete the old **.spec.metrics** field.

5. Update the **apiVersion** of the **KafkaConnect** custom resource to **v1beta2**:
   Replace:

```
apiVersion: kafka.strimzi.io/v1beta1
```

   with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

6. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.12. Upgrading Kafka Connect S2I resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **KafkaConnectS2I** custom resource in your deployment.

1. Update the **KafkaConnectS2I** custom resource in an editor.

```
oc edit kafkaconnects2i S2I-CLUSTER
```

2. If present, move:

```
KafkaConnectS2I.spec.affinity
```

```
KafkaConnectS2I.spec.tolerations
```

to:

```
KafkaConnectS2I.spec.template.pod.affinity
```

```
KafkaConnectS2I.spec.template.pod.tolerations
```

For example, move:

```
spec:
  # ...
  affinity:
    # ...
  tolerations:
    # ...
```

to:

```
spec:
  # ...
  template:
    pod:
      affinity:
        # ...
      tolerations:
        # ...
```

3. If **type: external** logging is configured in **.spec.logging**:
Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j.properties
```

4. If the **.spec.metrics** field is used to enable metrics:

   a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.metrics** field.

```
kind: ConfigMap
apiVersion: v1
metadata:
  name: kafka-connect-s2i-metrics
  labels:
    app: strimzi
data:
  connect-s2i-metrics-config.yaml: |
    <YAML>
```

b. Add a **.spec.metricsConfig** property that points to the ConfigMap and key:

```
metricsConfig:
  type: jmxPrometheusExporter
  valueFrom:
    configMapKeyRef:
      name: kafka-connect-s2i-metrics
      key: connect-s2i-metrics-config.yaml
```

c. Delete the old **.spec.metrics** field

5. Update the **apiVersion** of the **KafkaConnectS2I** custom resource to **v1beta2**: Replace:

```
apiVersion: kafka.strimzi.io/v1beta1
```

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

6. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.13. Upgrading Kafka MirrorMaker resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- MirrorMaker is configured and deployed. See Section 5.3.1, "Deploying Kafka MirrorMaker to your OpenShift cluster".

Procedure

Perform the following steps for each **KafkaMirrorMaker** custom resource in your deployment.

1. Update the **KafkaMirrorMaker** custom resource in an editor.

```
oc edit kafkamirrormaker MIRROR-MAKER
```

2. If present, move:

```
KafkaMirrorMaker.spec.affinity
```

```
KafkaMirrorMaker.spec.tolerations
```

to:

```
KafkaMirrorMaker.spec.template.pod.affinity
```

```
KafkaMirrorMaker.spec.template.pod.tolerations
```

For example, move:

```
spec:
  # ...
  affinity:
    # ...
  tolerations:
    # ...
```

to:

```
spec:
  # ...
  template:
    pod:
      affinity:
        # ...
      tolerations:
        # ...
```

3. If **type: external** logging is configured in **.spec.logging**:
   Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

   With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j.properties
```

4. If the **.spec.metrics** field is used to enable metrics:

   a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.metrics** field.

```
kind: ConfigMap
apiVersion: v1
metadata:
```

```
  name: kafka-mm-metrics
  labels:
    app: strimzi
data:
  mm-metrics-config.yaml: |
    <YAML>
```

b. Add a **.spec.metricsConfig** property that points to the ConfigMap and key:

```
metricsConfig:
  type: jmxPrometheusExporter
  valueFrom:
    configMapKeyRef:
      name: kafka-mm-metrics
      key: mm-metrics-config.yaml
```

c. Delete the old **.spec.metrics** field.

5. Update the **apiVersion** of the **KafkaMirrorMaker** custom resource to **v1beta2**:
Replace:

```
apiVersion: kafka.strimzi.io/v1beta1
```

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

6. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.14. Upgrading Kafka MirrorMaker 2.0 resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- MirrorMaker 2.0 is configured and deployed. See Section 5.3.1, “Deploying Kafka MirrorMaker to your OpenShift cluster”.

Procedure

Perform the following steps for each **KafkaMirrorMaker2** custom resource in your deployment.

1. Update the **KafkaMirrorMaker2** custom resource in an editor.

```
oc edit kafkamirrormaker2 MIRROR-MAKER-2
```

2. If present, move **affinity** from **.spec.affinity** to **.spec.template.pod.affinity**.
3. If present, move **tolerations** from **.spec.tolerations** to **.spec.template.pod.tolerations**.
4. If **type: external** logging is configured in **.spec.logging**:
Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j.properties
```

5. If the **.spec.metrics** field is used to enable metrics:

   a. Create a new ConfigMap that stores the YAML configuration for the JMX Prometheus exporter under a key. The YAML must match what is currently in the **.spec.metrics** field.

   ```
   kind: ConfigMap
   apiVersion: v1
   metadata:
     name: kafka-mm2-metrics
     labels:
       app: strimzi
   data:
     mm2-metrics-config.yaml: |
       <YAML>
   ```

   b. Add a **.spec.metricsConfig** property that points to the ConfigMap and key:

   ```
   metricsConfig:
     type: jmxPrometheusExporter
     valueFrom:
       configMapKeyRef:
         name: kafka-mm2-metrics
         key: mm2-metrics-config.yaml
   ```

   c. Delete the old **.spec.metrics** field.

6. Update the **apiVersion** of the **KafkaMirrorMaker2** custom resource to **v1beta2**: Replace:

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

   with:

   ```
   apiVersion: kafka.strimzi.io/v1beta2
   ```

7. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.15. Upgrading Kafka Bridge resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- The Kafka Bridge is configured and deployed. See Section 5.4.1, "Deploying Kafka Bridge to your OpenShift cluster".

Procedure

Perform the following steps for each **KafkaBridge** resource in your deployment.

1. Update the **KafkaBridge** custom resource in an editor.

```
oc edit kafkabridge KAFKA-BRIDGE
```

2. If **type: external** logging is configured in **KafkaBridge.spec.logging**:
   Replace the **name** of the ConfigMap containing the logging configuration:

```
logging:
  type: external
  name: my-config-map
```

   With the **valueFrom.configMapKeyRef** field, and specify both the ConfigMap **name** and the **key** under which the logging is stored:

```
logging:
  type: external
  valueFrom:
    configMapKeyRef:
      name: my-config-map
      key: log4j2.properties
```

3. Update the **apiVersion** of the **KafkaBridge** custom resource to **v1beta2**:
   Replace:

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

   with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

4. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.16. Upgrading Kafka User resources to v1beta2

Prerequisites

- A User Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **KafkaUser** custom resource in your deployment.

1. Update the **KafkaUser** custom resource in an editor.

```
oc edit kafkauser KAFKA-USER
```

2. Update the **apiVersion** of the **KafkaUser** custom resource to **v1beta2**:
Replace:

```
apiVersion: kafka.strimzi.io/v1beta1
```

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

3. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.17. Upgrading Kafka Topic resources to v1beta2

Prerequisites

- A Topic Operator supporting the **v1beta2** API version is up and running.

Procedure

Perform the following steps for each **KafkaTopic** custom resource in your deployment.

1. Update the **KafkaTopic** custom resource in an editor.

```
oc edit kafkatopic KAFKA-TOPIC
```

2. Update the **apiVersion** of the **KafkaTopic** custom resource to **v1beta2**:
Replace:

```
apiVersion: kafka.strimzi.io/v1beta1
```

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

3. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.18. Upgrading Kafka Connector resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- **KafkaConnector** custom resources are deployed to manage connector instances. See Section 5.2.4, “Creating and managing connectors”.

Procedure

Perform the following steps for each **KafkaConnector** custom resource in your deployment.

1. Update the **KafkaConnector** custom resource in an editor.

```
oc edit kafkaconnector KAFKA-CONNECTOR
```

2. Update the **apiVersion** of the **KafkaConnector** custom resource to **v1beta2**:
Replace:

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

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

3. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.2.19. Upgrading Kafka Rebalance resources to v1beta2

Prerequisites

- A Cluster Operator supporting the **v1beta2** API version is up and running.
- Cruise Control is configured and deployed. See Deploying Cruise Control in the *Using AMQ Streams on OpenShift* guide.

Procedure

Perform the following steps for each **KafkaRebalance** custom resource in your deployment.

1. Update the **KafkaRebalance** custom resource in an editor.

```
oc edit kafkarebalance KAFKA-REBALANCE
```

2. Update the **apiVersion** of the **KafkaRebalance** custom resource to **v1beta2**:
Replace:

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

with:

```
apiVersion: kafka.strimzi.io/v1beta2
```

3. Save the file, exit the editor and wait for the updated custom resource to be reconciled.

8.3. UPGRADING CONSUMERS TO COOPERATIVE REBALANCING

You can upgrade Kafka consumers and Kafka Streams applications to use the *incremental cooperative rebalance* protocol for partition rebalances instead of the default *eager rebalance* protocol. The new protocol was added in Kafka 2.4.0.

Consumers keep their partition assignments in a cooperative rebalance and only revoke them at the end of the process, if needed to achieve a balanced cluster. This reduces the unavailability of the consumer group or Kafka Streams application.

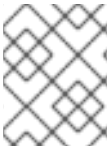

NOTE

Upgrading to the incremental cooperative rebalance protocol is optional. The eager rebalance protocol is still supported.

Prerequisites

- You have upgraded Kafka brokers and client applications to Kafka 2.7.0.

Procedure

To upgrade a Kafka consumer to use the incremental cooperative rebalance protocol:

1. Replace the Kafka clients **.jar** file with the new version.
2. In the consumer configuration, append **cooperative-sticky** to the **partition.assignment.strategy**. For example, if the **range** strategy is set, change the configuration to **range, cooperative-sticky**.
3. Restart each consumer in the group in turn, waiting for the consumer to rejoin the group after each restart.
4. Reconfigure each consumer in the group by removing the earlier **partition.assignment.strategy** from the consumer configuration, leaving only the **cooperative-sticky** strategy.
5. Restart each consumer in the group in turn, waiting for the consumer to rejoin the group after each restart.

To upgrade a Kafka Streams application to use the incremental cooperative rebalance protocol:

1. Replace the Kafka Streams **.jar** file with the new version.
2. In the Kafka Streams configuration, set the **upgrade.from** configuration parameter to the Kafka version you are upgrading from (for example, 2.3).
3. Restart each of the stream processors (nodes) in turn.
4. Remove the **upgrade.from** configuration parameter from the Kafka Streams configuration.
5. Restart each consumer in the group in turn.

Additional resources

- Notable changes in 2.4.0 in the Apache Kafka documentation.

CHAPTER 9. DOWNGRADING AMQ STREAMS

If you are encountering issues with the version of AMQ Streams you upgraded to, you can revert your installation to the previous version.

You can perform a downgrade to:

1. Revert your Cluster Operator to the previous AMQ Streams version.
   - Section 9.1, “Downgrading the Cluster Operator to a previous version”
2. Downgrade all Kafka brokers and client applications to the previous Kafka version.
   - Section 9.2, “Downgrading Kafka”

If the previous version of AMQ Streams does not support the version of Kafka you are using, you can also downgrade Kafka as long as the log message format versions appended to messages match.

9.1. DOWNGRADING THE CLUSTER OPERATOR TO A PREVIOUS VERSION

If you are encountering issues with AMQ Streams, you can revert your installation.

This procedure describes how to downgrade a Cluster Operator deployment to a previous version.

Prerequisites

- An existing Cluster Operator deployment is available.
- You have downloaded the installation files for the previous version .

Procedure

1. Take note of any configuration changes made to the existing Cluster Operator resources (in the **/install/cluster-operator** directory). Any changes will be **overwritten** by the previous version of the Cluster Operator.
2. Revert your custom resources to reflect the supported configuration options available for the version of AMQ Streams you are downgrading to.
3. Update the Cluster Operator.
   a. Modify the installation files for the previous version according to the namespace the Cluster Operator is running in.
   On Linux, use:

   ```
   sed -i 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

   On MacOS, use:

   ```
   sed -i '' 's/namespace: .*/namespace: my-cluster-operator-namespace/' install/cluster-operator/*RoleBinding*.yaml
   ```

b. If you modified one or more environment variables in your existing Cluster Operator **Deployment**, edit the **install/cluster-operator/060-Deployment-strimzi-cluster-operator.yaml** file to use those environment variables.

4. When you have an updated configuration, deploy it along with the rest of the installation resources:

```
oc replace -f install/cluster-operator
```

Wait for the rolling updates to complete.

5. Get the image for the Kafka pod to ensure the downgrade was successful:

```
oc get pod my-cluster-kafka-0 -o jsonpath='{.spec.containers[0].image}'
```

The image tag shows the new AMQ Streams version followed by the Kafka version. For example, ***NEW-STRIMZI-VERSION*-kafka-*CURRENT-KAFKA-VERSION***.

Your Cluster Operator was downgraded to the previous version.

9.2. DOWNGRADING KAFKA

Kafka version downgrades are performed by the Cluster Operator.

9.2.1. Kafka version compatibility for downgrades

Kafka downgrades are dependent on compatible current and target Kafka versions, and the state at which messages have been logged.

You cannot revert to the previous Kafka version if that version does not support any of the **inter.broker.protocol.version** settings which have *ever been used* in that cluster, or messages have been added to message logs that use a newer **log.message.format.version**.

The **inter.broker.protocol.version** determines the schemas used for persistent metadata stored by the broker, such as the schema for messages written to **__consumer_offsets**. If you downgrade to a version of Kafka that does not understand an **inter.broker.protocol.version** that has (ever) been previously used in the cluster the broker will encounter data it cannot understand.

If the target downgrade version of Kafka has:

- The *same* **log.message.format.version** as the current version, the Cluster Operator downgrades by performing a single rolling restart of the brokers.
- A *different* **log.message.format.version**, downgrading is only possible if the running cluster has *always* had **log.message.format.version** set to the version used by the downgraded version. This is typically only the case if the upgrade procedure was aborted before the **log.message.format.version** was changed. In this case, the downgrade requires:
  - Two rolling restarts of the brokers if the interbroker protocol of the two versions is different
  - A single rolling restart if they are the same

Downgrading is *not possible* if the new version has ever used a **log.message.format.version** that is not supported by the previous version, including when the default value for **log.message.format.version** is used. For example, this resource can be downgraded to Kafka version 2.6.0 because the **log.message.format.version** has not been changed:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
spec:
  # ...
  kafka:
    version: 2.7.0
    config:
      log.message.format.version: "2.6"
      # ...
```

The downgrade would not be possible if the **log.message.format.version** was set at **"2.7"** or a value was absent (so that the parameter took the default value for a 2.7.0 broker of 2.7).

9.2.2. Downgrading Kafka brokers and client applications

This procedure describes how you can downgrade a AMQ Streams Kafka cluster to a lower (previous) version of Kafka, such as downgrading from 2.7.0 to 2.6.0.

Prerequisites

For the **Kafka** resource to be downgraded, check:

- IMPORTANT: Compatibility of Kafka versions.
- The Cluster Operator, which supports both versions of Kafka, is up and running.
- The **Kafka.spec.kafka.config** does not contain options that are not supported by the Kafka version being downgraded to.
- The **Kafka.spec.kafka.config** has a **log.message.format.version** and **inter.broker.protocol.version** that is supported by the Kafka version being downgraded to.

Procedure

1. Update the Kafka cluster configuration.

```
oc edit kafka KAFKA-CONFIGURATION-FILE
```

2. Change the **Kafka.spec.kafka.version** to specify the previous version.
   For example, if downgrading from Kafka 2.7.0 to 2.6.0:

```
apiVersion: kafka.strimzi.io/v1beta2
kind: Kafka
spec:
  # ...
  kafka:
    version: 2.6.0 (1)
    config:
      log.message.format.version: "2.6" (2)
      inter.broker.protocol.version: "2.6" (3)
      # ...
```

(1) Kafka version is changed to the previous version.

2 Message format version is unchanged.

3 Inter-broker protocol version is unchanged.

NOTE

You must format the value of **log.message.format.version** and **inter.broker.protocol.version** as a string to prevent it from being interpreted as a floating point number.

3. If the image for the Kafka version is different from the image defined in **STRIMZI_KAFKA_IMAGES** for the Cluster Operator, update **Kafka.spec.kafka.image**. See Section 8.1.3.1, "Kafka version and image mappings"

4. Save and exit the editor, then wait for rolling updates to complete.
Check the update in the logs or by watching the pod state transitions:

```
oc logs -f CLUSTER-OPERATOR-POD-NAME | grep -E "Kafka version downgrade from [0-9.]+ to [0-9.]+, phase ([0-9]+) of \1 completed"
```

```
oc get pod -w
```

Check the Cluster Operator logs for an **INFO** level message:

```
Reconciliation #NUM(watch) Kafka(NAMESPACE/NAME): Kafka version downgrade from FROM-VERSION to TO-VERSION, phase 1 of 1 completed
```

5. Downgrade all client applications (consumers) to use the previous version of the client binaries. The Kafka cluster and clients are now using the previous Kafka version.

6. If you are reverting back to a version of AMQ Streams earlier than 0.22, which uses ZooKeeper for the storage of topic metadata, delete the internal topic store topics from the Kafka cluster.

```
oc run kafka-admin -ti --image=registry.redhat.io/amq7/amq-streams-kafka-27-rhel7:1.7.0 --rm=true --restart=Never -- ./bin/kafka-topics.sh --bootstrap-server localhost:9092 --topic __strimzi-topic-operator-kstreams-topic-store-changelog --delete && ./bin/kafka-topics.sh --bootstrap-server localhost:9092 --topic __strimzi_store_topic --delete
```

Additional resources

- Topic Operator topic store

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 **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 2021-06-04 11:11:03 UTC