Developer Guide

Amazon MQ

Amazon MQ: Developer Guide

Table of Contents

What is Amazon MQ?

Amazon MQ is a managed message broker service that makes it easy to migrate to a message broker in the cloud. A *message broker* allows software applications and components to communicate using various programming languages, operating systems, and formal messaging protocols. Currently, Amazon MQ supports Apache ActiveMQ Classic and RabbitMQ engine types.

Amazon MQ works with your existing applications and services without the need to manage, operate, or maintain your own messaging system.

Topics

How is Amazon MQ different from Amazon SQS or Amazon SNS?

Amazon MQ is a managed message broker service that provides compatibility with many popular message brokers. We recommend Amazon MQ for migrating applications from existing message brokers that rely on compatibility with APIs such as JMS or protocols such as AMQP 0-9-1, AMQP 1.0, MQTT, OpenWire, and STOMP.

Amazon SQS and Amazon SNS are queue and topic services that are highly scalable, simple to use, and don't require you to set up message brokers. We recommend these services for new applications that can benefit from nearly unlimited scalability and simple APIs.

How can I get started with Amazon MQ?

- To create your first broker with Amazon MQ, see Getting Started with Amazon MQ.
- To find out the guidelines and caveats that will help you make the most of Amazon MQ, see Working with Amazon MQ for ActiveMQ and Working with Amazon MQ for RabbitMQ
- To learn about Amazon MQ REST APIs, see the *Amazon MQ REST API Reference*.
- To learn about Amazon MQ AWS CLI commands, see Amazon MQ in the *AWS CLI Command Reference*.

We want to hear from you

We welcome your feedback. To contact us, visit the Amazon MQ Discussion Forum.

Setting up Amazon MQ

Before you can use Amazon MQ, you must complete the following steps.

Topics

Step 1: Prerequisites

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open https://portal.aws.amazon.com/billing/signup.
2. Follow the online instructions.

   Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

   When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to an administrative user, and use only the root user to perform tasks that require root user access.

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to https://aws.amazon.com/ and choosing **My Account**.

Create an administrative user

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the AWS Management Console as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

   For help signing in by using root user, see Signing in as the root user in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

   For instructions, see Enable a virtual MFA device for your AWS account root user (console) in the *IAM User Guide*.

Create an administrative user

1. Enable IAM Identity Center.

   For instructions, see Enabling AWS IAM Identity Center in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to an administrative user.

   For a tutorial about using the IAM Identity Center directory as your identity source, see Configure user access with the default IAM Identity Center directory in the *AWS IAM Identity Center User Guide*.

Sign in as the administrative user

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

  For help signing in using an IAM Identity Center user, see Signing in to the AWS access portal in the *AWS Sign-In User Guide*.

Step 2: create a user and get your AWS credentials

Users need programmatic access if they want to interact with AWS outside of the AWS Management Console. The way to grant programmatic access depends on the type of user that's accessing AWS.

To grant users programmatic access, choose one of the following options.

Which user needs programmatic access?	To	By
Workforce identity (Users managed in IAM Identity Center)	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Configuring the AWS CLI to use AWS IAM Identity Center in the *AWS Command Line Interface User Guide*. • For AWS SDKs, tools, and AWS APIs, see IAM Identity Center authentication in the *AWS SDKs and Tools Reference Guide*.
IAM	Use temporary credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions in Using temporary credentials with AWS resources in the *IAM User Guide*.
IAM	(Not recommended) Use long-term credentials to sign programmatic requests to the AWS CLI, AWS SDKs, or AWS APIs.	Following the instructions for the interface that you want to use. • For the AWS CLI, see Authenticating using IAM user credentials in the *AWS Command Line Interface User Guide*. • For AWS SDKs and tools, see Authenticate using long-term credentials in

Which user needs programmatic access?	To	By
		the *AWS SDKs and Tools Reference Guide*. • For AWS APIs, see Managing access keys for IAM users in the *IAM User Guide*.

Step 3: get ready to use the example codes

The following tutorials show how you can work with Amazon MQ brokers using the AWS Management Console as well as how to connect to your Amazon MQ for ActiveMQ and Amazon MQ for RabbitMQ brokers programatically. To use the ActiveMQ Java example code, you must install the Java Standard Edition Development Kit and make some changes to the code.

You can also create and manage brokers programmatically using Amazon MQ REST API and AWS SDKs.

Next steps

Now that you're prepared to work with Amazon MQ, get started by creating a broker. Depending on your broker engine type, you can then connect a Java application to your Amazon MQ for ActiveMQ broker or use the RabbitMQ Java client library to connect a JVM-based application to your Amazon MQ for RabbitMQ broker.

Getting started with Amazon MQ

This section will help you become more familiar with Amazon MQ by showing you how to create an Amazon MQ for ActiveMQ or RabbitMQ broker and how to connect your application to it.

Creating and connecting to a broker instance is slightly different for each broker engine. Choose one of the following engine types that you want to use for detailed information on creating and connecting to a broker. After you have created and connected to your broker, you can find instructions to help you delete it.

Topics

Prerequisites

Before you begin, complete the steps in Setting Up Amazon MQ.

Creating and connecting to an ActiveMQ broker

A *broker* is a message broker environment running on Amazon MQ. It is the basic building block of Amazon MQ. The combined description of the broker instance *class* (`m5`, `t3`) and *size* (`large`, `micro`) is a *broker instance type* (for example, `mq.m5.large`). For more information, see Broker.

Topics

Step 1: create an ActiveMQ broker

The first and most common Amazon MQ task is creating a broker. The following example shows how you can use the AWS Management Console to create a basic broker.

1. Sign in to the Amazon MQ console.
2. On the **Select broker engine** page, choose **Apache ActiveMQ**.
3. On the **Select deployment and storage** page, in the **Deployment mode and storage type** section, do the following:

   a. Choose the **Deployment mode** (for example, **Active/standby broker**). For more information, see Broker Architecture.

      - A **Single-instance broker** is comprised of one broker in one Availability Zone. The broker communicates with your application and with an Amazon EBS or Amazon EFS storage volume. For more information, see Amazon MQ single-instance broker.
      - An **Active/standby broker for high availability** is comprised of two brokers in two different Availability Zones, configured in a *redundant pair*. These brokers communicate synchronously with your application, and with Amazon EFS. For more information, see Amazon MQ active/standby broker for high availability.
      - For more information on the sample blueprints for a network of brokers, see Sample blueprints.

   b. Choose the **Storage type** (for example, **EBS**). For more information, see Storage.

      > **Note**
      >
      > Amazon EBS replicates data within a single Availability Zone and doesn't support the ActiveMQ active/standby deployment mode.

   c. Choose **Next**.
4. On the **Configure settings** page, in the **Details** section, do the following:

   a. Enter the **Broker name**.

      > **Important**
      >
      > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker names. Broker names are accessible to other AWS

services, including CloudWatch Logs. Broker names are not intended to be used for private or sensitive data.

b. Choose the **Broker instance type** (for example, **mq.m5.large**). For more information, see Broker instance types.

5. In the **ActiveMQ Web Console access** section, provide a **Username** and **Password**. The following restrictions apply to broker usernames and passwords:

   - Your username can contain only alphanumeric characters, dashes, periods, underscores, and tildas (- . _ ~).
   - Your password must be at least 12 characters long, contain at least 4 unique characters and must not contain commas, colons, or equal signs (,:=).

   > **Important**
   >
   > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

6. Choose **Deploy**.

   While Amazon MQ creates your broker, it displays the **Creation in progress** status.

   Creating the broker takes about 15 minutes.

   When your broker is created successfully, Amazon MQ displays the **Running** status.

	Name	Status	Deployment mode	Instance type
○	MyBroker	Running	Single-instance broker	mq.m5.large

7. Choose ***MyBroker***.

   On the ***MyBroker*** page, in the **Connect** section, note your broker's **ActiveMQ web console** URL, for example:

```
https://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:8162
```

Also, note your broker's wire-level protocol **Endpoints**. The following is an example of an OpenWire endpoint:

```
ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61617
```

Step 2: connect a Java application to your broker

After you create an Amazon MQ ActiveMQ broker, you can connect your application to it. The following examples show how you can use the Java Message Service (JMS) to create a connection to the broker, create a queue, and send a message. For a complete, working Java example, see Working Java Example.

You can connect to ActiveMQ brokers using various ActiveMQ clients. We recommend using the ActiveMQ Client.

Prerequisites

Enable VPC attributes

> **Note**
>
> You cannot disable public accessibility for your existing Amazon MQ brokers.

To ensure that your broker is accessible within your VPC, you must enable the `enableDnsHostnames` and `enableDnsSupport` VPC attributes. For more information, see DNS Support in your VPC in the *Amazon VPC User Guide*.

Enable inbound connections

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**).
3. On the ***MyBroker*** page, in the **Connections** section, note the addresses and ports of the broker's web console URL and wire-level protocols.

4. In the **Details** section, under **Security and network**, choose the name of your security group or .

   The **Security Groups** page of the EC2 Dashboard is displayed.
5. From the security group list, choose your security group.
6. At the bottom of the page, choose **Inbound**, and then choose **Edit**.
7. In the **Edit inbound rules** dialog box, add a rule for every URL or endpoint that you want to be publicly accessible (the following example shows how to do this for a broker web console).

   a. Choose **Add Rule**.

   b. For **Type**, select **Custom TCP**.

   c. For **Port Range**, type the web console port (`8162`).

   d. For **Source**, leave **Custom** selected and then type the IP address of the system that you want to be able to access the web console (for example, `192.0.2.1`).

   e. Choose **Save**.

   Your broker can now accept inbound connections.

Add Java dependencies

Add the `activemq-client.jar` and `activemq-pool.jar` packages to your Java class path. The following example shows these dependencies in a Maven project `pom.xml` file.

```
<dependencies>
    <dependency>
        <groupId>org.apache.activemq</groupId>
        <artifactId>activemq-client</artifactId>
        <version>5.15.16</version>
    </dependency>
    <dependency>
        <groupId>org.apache.activemq</groupId>
        <artifactId>activemq-pool</artifactId>
        <version>5.15.16</version>
    </dependency>
</dependencies>
```

For more information about `activemq-client.jar`, see Initial Configuration in the Apache ActiveMQ documentation.

> **Important**
>
> In the following example code, producers and consumers run in a single thread. For production systems (or to test broker instance failover), make sure that your producers and consumers run on separate hosts or threads.

Create a message producer and send a message

1. Create a JMS pooled connection factory for the message producer using your broker's endpoint and then call the `createConnection` method against the factory.

> **Note**
>
> For an active/standby broker, Amazon MQ provides two ActiveMQ Web Console URLs, but only one URL is active at a time. Likewise, Amazon MQ provides two endpoints for each wire-level protocol, but only one endpoint is active in each pair at a time. The `-1` and `-2` suffixes denote a redundant pair. For more information, see Broker Architecture).
> For wire-level protocol endpoints, you can allow your application to connect to either endpoint by using the Failover Transport.

```
// Create a connection factory.
final ActiveMQConnectionFactory connectionFactory = new
 ActiveMQConnectionFactory(wireLevelEndpoint);

// Pass the sign-in credentials.
connectionFactory.setUserName(activeMqUsername);
connectionFactory.setPassword(activeMqPassword);

// Create a pooled connection factory.
final PooledConnectionFactory pooledConnectionFactory = new
 PooledConnectionFactory();
pooledConnectionFactory.setConnectionFactory(connectionFactory);
pooledConnectionFactory.setMaxConnections(10);

// Establish a connection for the producer.
final Connection producerConnection = pooledConnectionFactory.createConnection();
```

```
producerConnection.start();

// Close all connections in the pool.
pooledConnectionFactory.clear();
```

> **Note**
>
> Message producers should always use the `PooledConnectionFactory` class. For more information, see Always Use Connection Pooling.

2. Create a session, a queue named `MyQueue`, and a message producer.

```
// Create a session.
final Session producerSession = producerConnection.createSession(false,
 Session.AUTO_ACKNOWLEDGE);

// Create a queue named "MyQueue".
final Destination producerDestination = producerSession.createQueue("MyQueue");

// Create a producer from the session to the queue.
final MessageProducer producer =
 producerSession.createProducer(producerDestination);
producer.setDeliveryMode(DeliveryMode.NON_PERSISTENT);
```

3. Create the message string `"Hello from Amazon MQ!"` and then send the message.

```
// Create a message.
final String text = "Hello from Amazon MQ!";
TextMessage producerMessage = producerSession.createTextMessage(text);

// Send the message.
producer.send(producerMessage);
System.out.println("Message sent.");
```

4. Clean up the producer.

```
producer.close();
producerSession.close();
producerConnection.close();
```

Create a message consumer and receive the message

1. Create a JMS connection factory for the message producer using your broker's endpoint and then call the `createConnection` method against the factory.

```
// Create a connection factory.
final ActiveMQConnectionFactory connectionFactory = new
 ActiveMQConnectionFactory(wireLevelEndpoint);

// Pass the sign-in credentials.
connectionFactory.setUserName(activeMqUsername);
connectionFactory.setPassword(activeMqPassword);

// Establish a connection for the consumer.
final Connection consumerConnection = connectionFactory.createConnection();
consumerConnection.start();
```

> **Note**
>
> Message consumers should *never* use the `PooledConnectionFactory` class. For more information, see Always Use Connection Pooling.

2. Create a session, a queue named `MyQueue`, and a message consumer.

```
// Create a session.
final Session consumerSession = consumerConnection.createSession(false,
 Session.AUTO_ACKNOWLEDGE);

// Create a queue named "MyQueue".
final Destination consumerDestination = consumerSession.createQueue("MyQueue");

// Create a message consumer from the session to the queue.
final MessageConsumer consumer =
 consumerSession.createConsumer(consumerDestination);
```

3. Begin to wait for messages and receive the message when it arrives.

```
// Begin to wait for messages.
final Message consumerMessage = consumer.receive(1000);

// Receive the message when it arrives.
```

```
final TextMessage consumerTextMessage = (TextMessage) consumerMessage;
System.out.println("Message received: " + consumerTextMessage.getText());
```

> **Note**
>
> Unlike AWS messaging services (such as Amazon SQS), the consumer is constantly connected to the broker.

4. Close the consumer, session, and connection.

```
consumer.close();
consumerSession.close();
consumerConnection.close();
```

Step 3: (Optional) connect to an AWS Lambda function

AWS Lambda can connect to and consume messages from your Amazon MQ broker. When you connect a broker to Lambda, you create an event source mapping that reads messages from a queue and invokes the function synchronously. The event source mapping you create reads messages from your broker in batches and converts them into a Lambda payload in the form of a JSON object.

To connect your broker to a Lambda function

1. Add the following IAM role permissions to your Lambda function execution role.

   - mq:DescribeBroker
   - ec2:CreateNetworkInterface
   - ec2:DeleteNetworkInterface
   - ec2:DescribeNetworkInterfaces
   - ec2:DescribeSecurityGroups
   - ec2:DescribeSubnets
   - ec2:DescribeVpcs
   - logs:CreateLogGroup
   - logs:CreateLogStream
   - logs:PutLogEvents

- secretsmanager:GetSecretValue

> **Note**
>
> Without the necessary IAM permissions, your function will not be able to successfully read records from Amazon MQ resources.

2. (Optional) If you have created a broker without public accessibility, you must do one of the following to allow Lambda to connect to your broker:

   - Configure one NAT gateway per public subnet. For more information, see Internet and service access for VPC-connected functions in the *AWS Lambda Developer Guide*.
   - Create a connection between your Amazon Virtual Private Cloud (Amazon VPC) and Lambda using a VPC endpoint. Your Amazon VPC must also connect to AWS Security Token Service (AWS STS) and Secrets Manager endpoints. For more information, see Configuring interface VPC endpoints for Lambda in the *AWS Lambda Developer Guide*.

3. Configure your broker as an event source for a Lambda function using the AWS Management Console. You can also use the `create-event-source-mapping` AWS Command Line Interface command.

4. Write some code for your Lambda function to process the messages consumed from your broker. The Lambda payload that retrieved by your event source mapping depends on the engine type of the broker. The following is an example of a Lambda payload for an Amazon MQ for ActiveMQ queue.

> **Note**
>
> In the example, `testQueue` is the name of the queue.

```
{
  "eventSource": "aws:amq",
  "eventSourceArn": "arn:aws:mq:us-
west-2:112556298976:broker:test:b-9bcfa592-423a-4942-879d-eb284b418fc8",
  "messages": {
    [
      {
```

```
        "messageID": "ID:b-9bcfa592-423a-4942-879d-eb284b418fc8-1.mq.us-
west-2.amazonaws.com-37557-1234520418293-4:1:1:1:1",
        "messageType": "jms/text-message",
        "data": "QUJDOkFBQUE=",
        "connectionId": "myJMSCoID",
        "redelivered": false,
        "destination": {
          "physicalname": "testQueue"
        },
        "timestamp": 1598827811958,
        "brokerInTime": 1598827811958,
        "brokerOutTime": 1598827811959
      },
      {
        "messageID": "ID:b-9bcfa592-423a-4942-879d-eb284b418fc8-1.mq.us-
west-2.amazonaws.com-37557-1234520418293-4:1:1:1:1",
        "messageType":"jms/bytes-message",
        "data": "3DTOOW7crj51prgVLQaGQ82S48k=",
        "connectionId": "myJMSCoID1",
        "persistent": false,
        "destination": {
          "physicalname": "testQueue"
        },
        "timestamp": 1598827811958,
        "brokerInTime": 1598827811958,
        "brokerOutTime": 1598827811959
      }
    ]
  }
}
```

For more information about connecting Amazon MQ to Lambda, the options Lambda supports for an Amazon MQ event source, and event source mapping errors, see Using Lambda with Amazon MQ in the *AWS Lambda Developer Guide*.

Step 4: delete your broker

If you don't use an Amazon MQ broker (and don't foresee using it in the near future), it is a best practice to delete it from Amazon MQ to reduce your AWS costs.

The following example shows how you can delete a broker using the AWS Management Console.

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Delete**.
3. In the **Delete *MyBroker*?** dialog box, type `delete` and then choose **Delete**.

   Deleting a broker takes about 5 minutes.

Next steps

Now that you have created a broker, connected an application to it, and sent and received a message, you might want to try the following:

- Creating and configuring a broker (Additional Settings)
- Editing broker engine version, instance type, CloudWatch logs, and maintenance preferences
- Creating and applying broker configurations
- Listing brokers and viewing broker details
- Creating and managing ActiveMQ broker users
- Rebooting a Broker
- Accessing CloudWatch metrics for Amazon MQ

You can also begin to dive deep into Best Practices for Amazon MQ and Amazon MQ REST APIs, and then plan to migrate to Amazon MQ.

Creating and connecting to a RabbitMQ broker

A *broker* is a message broker environment running on Amazon MQ. It is the basic building block of Amazon MQ. The combined description of the broker instance *class* (`m5`, `t3`) and *size* (`large`, `micro`) is a *broker instance type* (for example, `mq.m5.large`).

Topics

- Step 1: create a RabbitMQ broker
- Step 2: connect a JVM-based application to your broker
- Step 3: (Optional) connect to an AWS Lambda function
- Step 4: delete your broker

- Next steps

Step 1: create a RabbitMQ broker

The first and most common Amazon MQ task is creating a broker. The following example shows how you can use the AWS Management Console to create a basic broker.

1. Sign in to the Amazon MQ console.
2. On the **Select broker engine** page, choose **RabbitMQ**, and then choose **Next**.
3. On the **Select deployment mode** page, choose the **Deployment mode**, for example, **Cluster deployment**, and then choose **Next**.

   - A **single-instance broker** is comprised of one broker in one Availability Zone behind a Network Load Balancer (NLB). The broker communicates with your application and with an Amazon EBS storage volume. For more information, see Single-instance broker.
   - A **RabbitMQ cluster deployment for high availability** is a logical grouping of three RabbitMQ broker nodes behind a Network Load Balancer, each sharing users, queues, and a distributed state across multiple Availability Zones (AZ). For more information, see Cluster deployment for high availability.
4. On the **Configure settings** page, in the **Details** section, the following:

   a. Enter the Broker name.

   > **Important**
   >
   > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker names. Broker names are accessible to other AWS services, including CloudWatch Logs. Broker names are not intended to be used for private or sensitive data.

   b. Choose the **Broker instance type** (for example, **mq.m5.large**). For more information, see Broker instance types.

   > **Note**
   >
   > The **Additional settings** section provides options to enable CloudWatch logs and configure network access for your broker. If you create a private RabbitMQ broker

without public accessibility, you must select a Virtual Private Cloud (VPC) and configure a security group to access your broker.

5. On the **Configure settings** page, in the **RabbitMQ access** section, provide a **Username** and **Password**. The following restrictions apply to broker sign-in credentials:

   - Your username can contain only alphanumeric characters, dashes, periods, and underscores (- . _). This value must not contain any tilde (~) characters. Amazon MQ prohibits using `guest` as a username.
   - Your password must be at least 12 characters long, contain at least 4 unique characters and must not contain commas, colons, or equal signs (,:=).

   > **Important**
   >
   > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

6. Choose **Next**.
7. On the **Review and create** page, you can review your selections and edit them as needed.
8. Choose **Create broker**.

   While Amazon MQ creates your broker, it displays the **Creation in progress** status.

   Creating the broker takes about 15 minutes.

   When your broker is created successfully, Amazon MQ displays the **Running** status.

	Name	Status	Deployment mode	Instance type
○	MyBroker	Running	Single-instance broker	mq.m5.large

9. Choose ***MyBroker***.

   On the ***MyBroker*** page, in the **Connect** section, note your broker's **RabbitMQ web console** URL, for example:

```
https://b-c8349341-ec91-4a78-ad9c-a57f23f235bb.mq.us-west-2.amazonaws.com
```

Also, note your broker's secure-AMQP **Endpoint**. The following is an example of an `amqps` endpoint exposing listener port 5671.

```
amqps://b-c8349341-ec91-4a78-ad9c-a57f23f235bb.mq.us-west-2.amazonaws.com:5671
```

Step 2: connect a JVM-based application to your broker

After you create a RabbitMQ broker, you can connect your application to it. The following examples show how you can use the RabbitMQ Java client library to create a connection to your broker, create a queue, and send a message. You can connect to RabbitMQ brokers using supported RabbitMQ client libraries for a variety of languages. For more information about supported RabbitMQ client libraries, see RabbitMQ client libraries and developer tools.

Prerequisites

> **Note**
>
> The following prerequisite steps are only applicable to RabbitMQ brokers created without public accessibility. If you are creating a broker with public accessibility you can skip them.

Enable VPC attributes

To ensure that your broker is accessible within your VPC, you must enable the `enableDnsHostnames` and `enableDnsSupport` VPC attributes. For more information, see DNS Support in your VPC in the *Amazon VPC User Guide*.

Enable inbound connections

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**).
3. On the ***MyBroker*** page, in the **Connections** section, note the addresses and ports of the broker's web console URL and wire-level protocols.

4. In the **Details** section, under **Security and network**, choose the name of your security group or .

   The **Security Groups** page of the EC2 Dashboard is displayed.
5. From the security group list, choose your security group.
6. At the bottom of the page, choose **Inbound**, and then choose **Edit**.
7. In the **Edit inbound rules** dialog box, add a rule for every URL or endpoint that you want to be publicly accessible (the following example shows how to do this for a broker web console).

   a. Choose **Add Rule**.

   b. For **Type**, select **Custom TCP**.

   c. For **Source**, leave **Custom** selected and then type the IP address of the system that you want to be able to access the web console (for example, `192.0.2.1`).

   d. Choose **Save**.

   Your broker can now accept inbound connections.

Add Java dependencies

If you are using Apache Maven for automating builds, add the following dependency to your `pom.xml` file. For more information about Project Object Model files in Apache Maven, see Introduction to the POM.

```
<dependency>
    <groupId>com.rabbitmq</groupId>
    <artifactId>amqp-client</artifactId>
    <version>5.9.0</version>
</dependency>
```

If you are using Gradle for automating builds, declare the following dependency.

```
dependencies {
    compile 'com.rabbitmq:amqp-client:5.9.0'
}
```

Import `Connection` and `Channel` classes

RabbitMQ Java client uses `com.rabbitmq.client` as its top-level package, with `Connection` and `Channel` API classes representing an AMQP 0-9-1 connection and channel, respectively. Import the `Connection` and `Channel` classes before using them, as shown in the following example.

```
import com.rabbitmq.client.Connection;
import com.rabbitmq.client.Channel;
```

Create a `ConnectionFactory` and connect to your broker

Use the following example to create an instance of the `ConnectionFactory` class with the given parameters. Use the `setHost` method to configure the broker endpoint you noted earlier. For AMQPS wire-level connections, use port 5671.

```
ConnectionFactory factory = new ConnectionFactory();

factory.setUsername(username);
factory.setPassword(password);

//Replace the URL with your information
factory.setHost("b-c8352341-ec91-4a78-ad9c-a43f23d325bb.mq.us-west-2.amazonaws.com");
factory.setPort(5671);

// Allows client to establish a connection over TLS
factory.useSslProtocol();

// Create a connection
Connection conn = factory.newConnection();

// Create a channel
Channel channel = conn.createChannel();
```

Publish a message to an exchange

You can use `Channel.basicPublish` to publish messages to an exchange. The following example uses the AMQP `Builder` class to build a message properties object with content-type `plain/text`.

```
byte[] messageBodyBytes = "Hello, world!".getBytes();
```

```
channel.basicPublish(exchangeName, routingKey,
             new AMQP.BasicProperties.Builder()
               .contentType("text/plain")
               .userId("userId")
               .build(),
               messageBodyBytes);
```

> **Note**
>
> Note that `BasicProperties` is an inner class of the autogenerated holder class, `AMQP`.

Subscribe to a queue and receive a message

You can receive a message by subscribing to a queue using the `Consumer` interface. Once subscribed, messages will then be delivered automatically as they arrive.

The easiest way to implement a `Consumer` is to use the subclass `DefaultConsumer`. A `DefaultConsumer` object can be passed as part of a `basicConsume` call to set up the subscription as shown in the following example.

```
boolean autoAck = false;
channel.basicConsume(queueName, autoAck, "myConsumerTag",
     new DefaultConsumer(channel) {
         @Override
         public void handleDelivery(String consumerTag,
                                    Envelope envelope,
                                    AMQP.BasicProperties properties,
                                    byte[] body)
             throws IOException
         {
             String routingKey = envelope.getRoutingKey();
             String contentType = properties.getContentType();
             long deliveryTag = envelope.getDeliveryTag();
             // (process the message components here ...)
             channel.basicAck(deliveryTag, false);
         }
     });
```

> **Note**
>
> Because we specified `autoAck = false`, it is necessary to acknowledge messages delivered to the `Consumer`, most conveniently done in the `handleDelivery` method, as shown in the example.

Close your connection and disconnect from the broker

In order to disconnect from your RabbitMQ broker, close both the channel and connection as shown in the following.

```
channel.close();
conn.close();
```

> **Note**
>
> For more information about working with the RabbitMQ Java client library, see the RabbitMQ Java Client API Guide.

Step 3: (Optional) connect to an AWS Lambda function

AWS Lambda can connect to and consume messages from your Amazon MQ broker. When you connect a broker to Lambda, you create an event source mapping that reads messages from a queue and invokes the function synchronously. The event source mapping you create reads messages from your broker in batches and converts them into a Lambda payload in the form of a JSON object.

To connect your broker to a Lambda function

1. Add the following IAM role permissions to your Lambda function execution role.

   - mq:DescribeBroker
   - ec2:CreateNetworkInterface
   - ec2:DeleteNetworkInterface
   - ec2:DescribeNetworkInterfaces
   - ec2:DescribeSecurityGroups

- ec2:DescribeSubnets
- ec2:DescribeVpcs
- logs:CreateLogGroup
- logs:CreateLogStream
- logs:PutLogEvents
- secretsmanager:GetSecretValue

> **Note**
>
> Without the necessary IAM permissions, your function will not be able to successfully read records from Amazon MQ resources.

2. (Optional) If you have created a broker without public accessibility, you must do one of the following to allow Lambda to connect to your broker:

   - Configure one NAT gateway per public subnet. For more information, see Internet and service access for VPC-connected functions in the *AWS Lambda Developer Guide*.
   - Create a connection between your Amazon Virtual Private Cloud (Amazon VPC) and Lambda using a VPC endpoint. Your Amazon VPC must also connect to AWS Security Token Service (AWS STS) and Secrets Manager endpoints. For more information, see Configuring interface VPC endpoints for Lambda in the *AWS Lambda Developer Guide*.

3. Configure your broker as an event source for a Lambda function using the AWS Management Console. You can also use the `create-event-source-mapping` AWS Command Line Interface command.

4. Write some code for your Lambda function to process the messages from your consumed from your broker. The Lambda payload that retrieved by your event source mapping depends on the engine type of the broker. The following is an example of a Lambda payload for an Amazon MQ for RabbitMQ queue.

> **Note**
>
> In the example, `test` is the name of the queue, and `/` is the name of the default virtual host. When receiving messages, the event source lists messages under `test::/`.

```
{
  "eventSource": "aws:rmq",
  "eventSourceArn": "arn:aws:mq:us-
west-2:112556298976:broker:test:b-9bcfa592-423a-4942-879d-eb284b418fc8",
  "rmqMessagesByQueue": {
    "test::/": [
      {
        "basicProperties": {
          "contentType": "text/plain",
          "contentEncoding": null,
          "headers": {
            "header1": {
              "bytes": [
                118,
                97,
                108,
                117,
                101,
                49
              ]
            },
            "header2": {
              "bytes": [
                118,
                97,
                108,
                117,
                101,
                50
              ]
            },
            "numberInHeader": 10
          }
          "deliveryMode": 1,
          "priority": 34,
          "correlationId": null,
          "replyTo": null,
          "expiration": "60000",
          "messageId": null,
          "timestamp": "Jan 1, 1970, 12:33:41 AM",
          "type": null,
          "userId": "AIDACKCEVSQ6C2EXAMPLE",
```

```
            "appId": null,
            "clusterId": null,
            "bodySize": 80
          },
          "redelivered": false,
          "data": "eyJ0aW1lb3V0IjowLCJkYXRhIjoiQ1pybWYwR3c4T3Y0YnFMUXhENEUifQ=="
        }
      ]
    }
  }
```

For more information about connecting Amazon MQ to Lambda, the options Lambda supports for an Amazon MQ event source, and event source mapping errors, see Using Lambda with Amazon MQ in the *AWS Lambda Developer Guide*.

Step 4: delete your broker

If you don't use an Amazon MQ broker (and don't foresee using it in the near future), it is a best practice to delete it from Amazon MQ to reduce your AWS costs.

The following example shows how you can delete a broker using the AWS Management Console.

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Delete**.
3. In the **Delete *MyBroker*?** dialog box, type `delete` and then choose **Delete**.

   Deleting a broker takes about 5 minutes.

Next steps

Now that you have created a broker, connected an application to it, and sent and received a message, you might want to try the following:

- Editing broker engine version, instance type, CloudWatch logs, and maintenance preferences
- Listing brokers and viewing broker details
- Creating and managing ActiveMQ broker users
- Rebooting a Broker
- Accessing CloudWatch metrics for Amazon MQ

You can also begin to dive deep into Best Practices for Amazon MQ and Amazon MQ REST APIs before planning to migrate to Amazon MQ.

Managing an Amazon MQ broker

In the following sections, you can find instructions for managing and maintaining your Amazon MQ brokers.

Topics

Maintaining an Amazon MQ broker

Periodically, Amazon MQ performs maintenance to the hardware, operating system, or the engine software a message broker. The duration of the maintenance varies, but can last up to two hours, depending on the operations that are scheduled for your message broker. For example, if you've activated automatic minor engine version upgrades, or changed the broker instance type, Amazon MQ will apply your changes during the next scheduled maintenance window.

To minimize downtime during a maintenance window, we recommend selecting a broker deployment mode with high availability across multiple Availability Zones (AZ). Depending on your broker engine type, Amazon MQ provides the following Multi-AZ deployment modes.

- **Amazon MQ for ActiveMQ** – Amazon MQ for ActiveMQ provides active/standby deployments for high availability. In active/standby mode, Amazon MQ performs maintenance operations one instance at a time, ensuring that at least one instance remains available. In addition, you can configure a network of brokers with maintenance windows scattered across the week.

- **Amazon MQ for RabbitMQ** – Amazon MQ for RabbitMQ provides the cluster deployments for high availability. In cluster deployments, Amazon MQ performs maintenance operations, one node at a time, keeping at least two running nodes at all times.

For more information about Amazon MQ recommended best practices to ensure your brokers perform effectively during, and after a maintenance window, see the following documentation for your broker engine type.

- the section called "Amazon MQ for ActiveMQ best practices"
- the section called "Amazon MQ for RabbitMQ best practices"

You can schedule maintenance to occur once a week at a specified time which lasts up to two hours. This sets the window for maintenance actions from Amazon MQ to be scheduled and started.

You can schedule the maintenance window when you first create your broker, or by updating your broker preferences. The following topic describes adjusting the broker maintenance window using the AWS Management Console, AWS CLI, and the Amazon MQ API.

Topics

Adjusting the broker maintenance window

During the maintenance window time you select, Amazon MQ will perform any pending changes, such as automatic minor version upgrades. To adjust the broker maintenance window, you can use the AWS Management Console, the AWS CLI, or the Amazon MQ API.

> **Important**
>
> You can only adjust the maintenance window of a broker up to **four** times before the next scheduled maintenance window. Amazon MQ applies a limit of four maintenance window adjustments to ensure that critical software and security patches, as well as important hardware upgrades, are not indefinitely deferred and postponed.
> Once a broker maintenance window is completed, Amazon MQ resets the limit, allowing you to adjust the schedule before the next maintenance window occurs.

Broker availability is not affected when adjusting the broker maintenance window.

AWS Management Console

To adjust the broker maintenance window by using the AWS Management Console

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**, and then choose the broker that you want to upgrade from the list.
3. On the broker details page, choose **Edit**.
4. Under **Maintenance**, do the following.
   a. For **Start day**, choose a day of the week, for example, **Sunday**, from the drop-down list.
   b. For **Start time**, choose the hour and minute of the day that you want to schedule for the next broker maintenance window, for example, **12:00**.

   > **Note**
   > The **Start time** options are configured in UTC+0 time zone.

5. Scroll to the bottom of the page, and choose **Save**. The maintenance window is adjusted immediately.
6. On the broker details page, under **Maintenance window**, verify that your new preferred schedule is displayed.

AWS CLI

To adjust the broker maintenance window using the AWS CLI

1. Use the update-broker CLI command and specify the following parameters, as shown in the example.
   - `--broker-id` – The unique ID that Amazon MQ generates for the broker. You can parse the ID from your broker ARN. For example, given the following ARN, `arn:aws:mq:us-east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`, the broker ID would be `b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`.

- `--maintenance-window-start-time` – The parameters that determine the weekly maintenance window start time provided in the following structure.
  - `DayOfWeek` – The day of the week, in the following syntax: `MONDAY| TUESDAY | WEDNESDAY | THURSDAY | FRIDAY | SATURDAY | SUNDAY`
  - `TimeOfDay` – The time, in 24-hour format.
  - `TimeZone` – (Optional) The time zone, in either the Country/City, or the UTC offset format. Set to UTC by default.

```
aws mq update-broker --broker-id broker-id \
--maintenance-window-start-time DayOfWeek=SUNDAY,TimeOfDay=13:00,TimeZone=America/
Los_Angeles
```

2. (Optional) Use the describe-broker CLI command to verify that the maintenance window is successfully updated.

```
aws mq describe-broker --broker-id broker-id
```

Amazon MQ API

To adjust the broker maintenance window using the Amazon MQ API

1. Use the UpdateBroker API operation. Specify `broker-id` as a path parameter. The following examples assumes a broker in the `us-west-2` region. For more information about available Amazon MQ endpoints, see Amazon MQ endpoints and quotas. in the *AWS General Reference*

```
PUT /v1/brokers/broker-id HTTP/1.1
Host: mq.us-west-2.amazonaws.com
Date: Wed, 7 July 2021 12:00:00 GMT
x-amz-date: Wed, 7 July 2021 12:00:00 GMT
Authorization: authorization-string
```

Use the `maintenanceWindowStartTime` parameter and the `WeeklyStartTime` resource type in the request payload.

```
{
"maintenanceWindowStartTime": {
    "dayOfWeek": "SUNDAY",
```

```
    "timeZone": "America/Los_Angeles",
    "timeOfDay": "13:00"
  }
}
```

2. (Optional) Use the DescribeBroker API operation to verify that the maintenance window has been successfully updated. `broker-id` is specified as a path parameter.

```
GET /v1/brokers/broker-id HTTP/1.1
Host: mq.us-west-2.amazonaws.com
Date: Wed, 7 July 2021 12:00:00 GMT
x-amz-date: Wed, 7 July 2021 12:00:00 GMT
Authorization: authorization-string
```

Upgrading an Amazon MQ broker engine version

Amazon MQ provides new broker engine versions for all supported broker engine type. New engine versions might include security patches, bug fixes and other broker engine improvements. When Amazon MQ supports a new engine version, you can control how and when to upgrade your broker.

Broker engine versions are organized as `X.Y.Z`. In the Amazon MQ implementation of each engine type, `X.Y` is considered a major version and `Z` is considered a minor version. There are two types of upgrades:

- **Major version upgrade** – Occurs when the major engine version numbers change. For example, upgrading from version 1.**0** to version 1.**1** is considered a major version upgrade.
- **Minor version upgrade** – Occurs when only the minor engine version number changes. For example, upgrading from version 1.1.**0** to version 1.1.**1** is considered a minor version upgrade.

For more information about major and minor version management for each specific broker engine type, see the following topics.

- the section called “Version management”
- the section called “Version management”

When you activate the automatic minor version upgrade option, Amazon MQ upgrades your broker to new minor versions as they become available. Automatic minor version upgrades occur only

if the broker is running a minor engine version that is lower than the new recommended minor version. For major upgrades, you must manually upgrade the engine version.

Both manual and automatic version upgrades occur during the scheduled maintenance window or after you reboot your broker.

The following topics describe how you can manually upgrade the broker engine version, and activate automatic minor version upgrades.

Topics

Manually upgrading the engine version

To manually upgrade the engine version of a broker to a new major or minor version, you can use the AWS Management Console, the AWS CLI, or the Amazon MQ API.

AWS Management Console

To upgrade the engine version of a broker by using the AWS Management Console

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**, and then choose the broker that you want to upgrade from the list.
3. On the broker details page, choose **Edit**.
4. Under **Specifications**, for **Broker engine version** choose the new version number from the dropdown list.
5. Scroll to the bottom of the page, and choose **Schedule modifications**.
6. On the **Schedule broker modifications** page, for **When to apply modifications**, choose one of the following.
   - Choose **After the next reboot**, if you want Amazon MQ to complete the version upgrade during the next scheduled maintenance window.
   - Choose **Immediately**, if you want to reboot the broker and upgrade the engine version immediately.

> **Important**
>
> Your broker will be offline while it is being rebooted.

7. Choose **Apply** to finish applying the changes.

AWS CLI

To upgrade the engine version of a broker by using the AWS CLI

1. Use the update-broker CLI command and specify the following parameters, as shown in the example.

   - `--broker-id` – The unique ID that Amazon MQ generates for the broker. You can parse the ID from your broker ARN. For example, given the following ARN, `arn:aws:mq:us-east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`, the broker ID would be `b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`.
   - `--engine-version` – The engine version number for the broker engine to upgrade to.

```
aws mq update-broker --broker-id broker-id --engine-version version-number
```

2. (Optional) Use the reboot-broker CLI command to reboot your broker if, you want to upgrade the engine version immediately.

```
aws mq reboot-broker --broker-id broker-id
```

   If you do not want to reboot your broker and apply the changes immediately, Amazon MQ will upgrade the broker during the next scheduled maintenance window.

> **Important**
>
> Your broker will be offline while it is being rebooted.

Amazon MQ API

To upgrade the engine version of a broker by using the Amazon MQ API

1. Use the UpdateBroker API operation. Specify `broker-id` as a path parameter. The following examples assumes a broker in the `us-west-2` region. For more information about available Amazon MQ endpoints, see Amazon MQ endpoints and quotas. in the *AWS General Reference*

```
PUT /v1/brokers/broker-id HTTP/1.1
Host: mq.us-west-2.amazonaws.com
Date: Mon, 7 June 2021 12:00:00 GMT
x-amz-date: Mon, 7 June 2021 12:00:00 GMT
Authorization: authorization-string
```

   Use `engineVersion` in the request payload to specify the version number for the broker to upgrade to.

```
{
    "engineVersion": "engine-version-number"
}
```

2. (Optional) Use the RebootBroker API operation to reboot your broker, if you want to upgrade the engine version immediately. `broker-id` is specified as a path parameter.

```
POST /v1/brokers/broker-id/reboot-broker HTTP/1.1
Host: mq.us-west-2.amazonaws.com
Date: Mon, 7 June 2021 12:00:00 GMT
x-amz-date: Mon, 7 June 2021 12:00:00 GMT
Authorization: authorization-string
```

   If you do not want to reboot your broker and apply the changes immediately, Amazon MQ will upgrade the broker during the next scheduled maintenance window.

   > **Important**
   >
   > Your broker will be offline while it is being rebooted.

Automatically upgrading the minor engine version

You can control whether automatic minor version upgrade is activated for a broker when you first create the broker, or by modifying broker preferences. To activate auto minor version upgrades for an existing broker, you can use the AWS Management Console, the AWS CLI, or the Amazon MQ API.

AWS Management Console

To activate automatic minor version upgrades by using the AWS Management Console

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**, and then choose the broker that you want to upgrade from the list.
3. On the broker details page, choose **Edit**.
4. Under **Maintenance**, choose **Enable automatic minor version upgrades**.

> **Note**
>
> If the option is already selected, you do not need to make any changes.

5. Choose **Save** at the bottom of the page.

AWS CLI

To activate automatic minor version upgrades via the AWS CLI, use the update-broker CLI command and specify the following parameters.

- `--broker-id` – The unique ID that Amazon MQ generates for the broker. You can parse the ID from your broker ARN. For example, given the following ARN, `arn:aws:mq:us-east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`, the broker ID would be `b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`.
- `--auto-minor-version-upgrade` – Activates the auto minor version upgrade option.

```
aws mq update-broker --broker-id broker-id --auto-minor-version-upgrade
```

If you want to deactivate auto minor version upgrades for your broker, use the `--no-auto-minor-version-upgrade` parameter.

Amazon MQ API

To activate automatic minor version upgrades via the Amazon MQ API, use the UpdateBroker API operation. Specify `broker-id` as a path parameter. The following example assumes a broker in the `us-west-2` region. For more information about available Amazon MQ endpoints, see Amazon MQ endpoints and quotas. in the *AWS General Reference*

```
PUT /v1/brokers/broker-id HTTP/1.1
Host: mq.us-west-2.amazonaws.com
Date: Mon, 7 June 2021 12:00:00 GMT
x-amz-date: Mon, 7 June 2021 12:00:00 GMT
Authorization: authorization-string
```

Use the `autoMinorVersionUpgrade` property in the request payload to activate auto minor version upgrade.

```
{
    "autoMinorVersionUpgrade": "true"
}
```

If you want to deactivate auto minor version upgrades for your broker, set `"autoMinorVersionUpgrade": "false"` in the request payload.

Broker statuses

A broker's current condition is indicated by a *status*. The following table lists the statuses of an Amazon MQ broker.

Console	API	Description
Creation failed	`CREATION_FAILED`	The broker couldn't be created.
Creation in progress	`CREATION_IN_PROGRESS`	The broker is currently being created.

Console	API	Description
Deletion in progress	`DELETION_IN_PROGRESS`	The broker is currently being deleted.
Reboot in progress	`REBOOT_IN_PROGRESS`	The broker is currently being rebooted.
Running	`RUNNING`	The broker is operational.
Critical action required	`CRITICAL_ACTION_REQUIRED`	The broker is running, but is in a degraded state and requires immediate action. You can find instructions to resolve the issue by chosing the action required code from the list in the section called "Troubleshooting: Amazon MQ action required codes".

Listing Amazon MQ brokers and viewing broker details

When you request that Amazon MQ create a broker, the creation process can take about 15 minutes.

The following example shows how you can confirm your broker's existence by listing your brokers in the current region using the AWS Management Console.

To list brokers and view broker details

1. Sign in to the Amazon MQ console.

   Your brokers in the current region are listed.

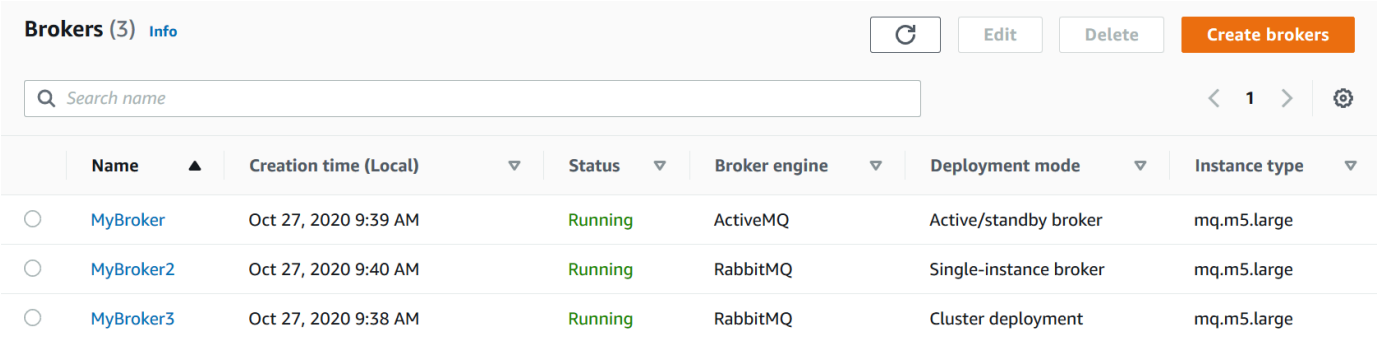

The following information is displayed for each broker:

- **Name**
- **Creation** date
- **Status**
- **Deployment mode**
- **Instance type**

2. Choose your broker's name .

For ActiveMQ brokers, on the ***MyBroker*** page, the configured **Details** are displayed for your broker:

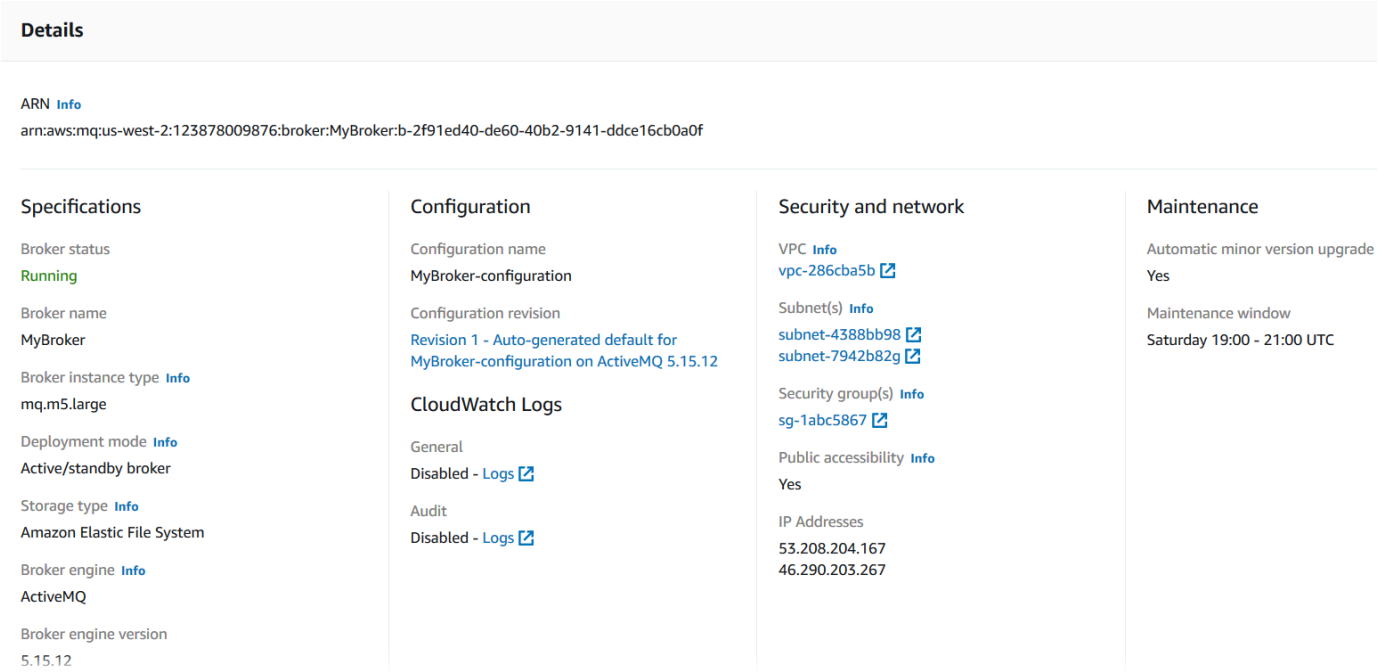

For Amazon MQ for RabbitMQ brokers, you can view your selected settings on the ***MyBroker2*** page, under the **Details**section as shown in the following.

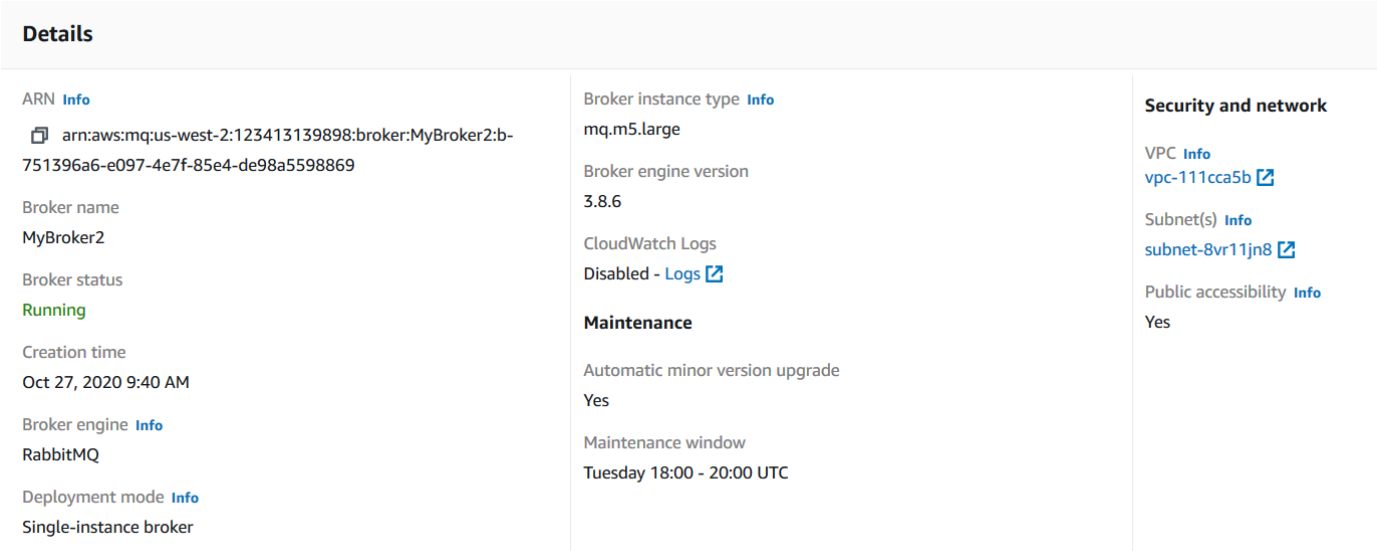

Under the **Details** section, the following information is displayed:

- In the **Connections** section, for Amazon MQ for ActiveMQ brokers, the web console URL and the wire-level protocol endpoints.

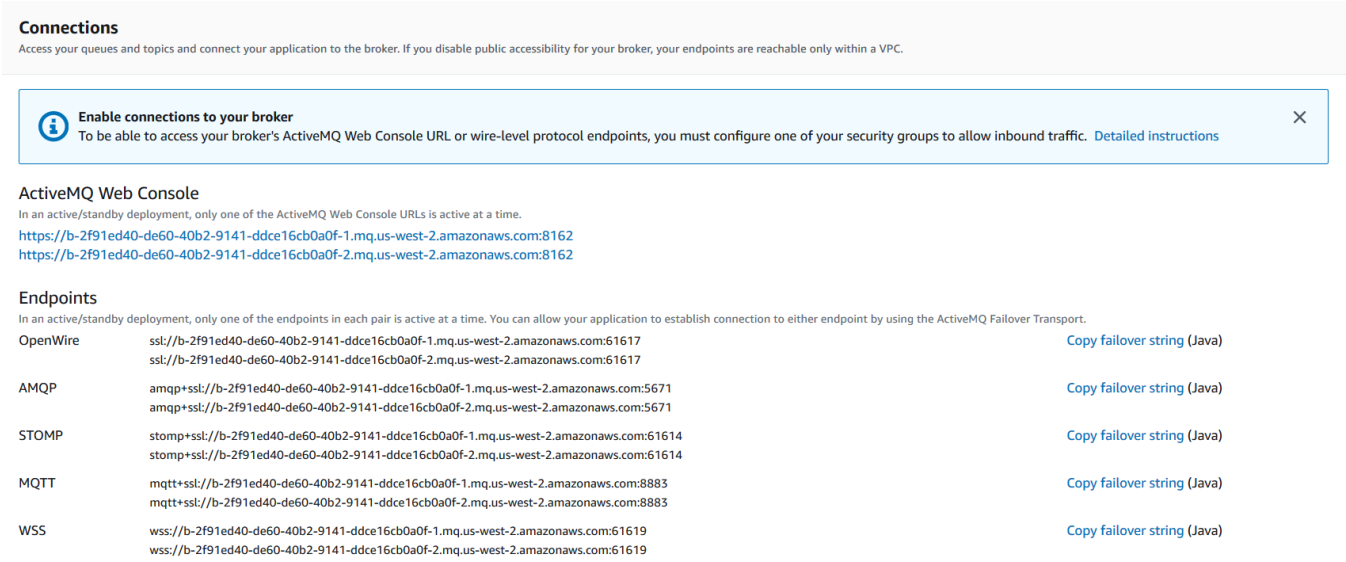

  In the **Connections** section, for Amazon MQ for RabbitMQ brokers, the web console URL and the secure AMQP endpoint.

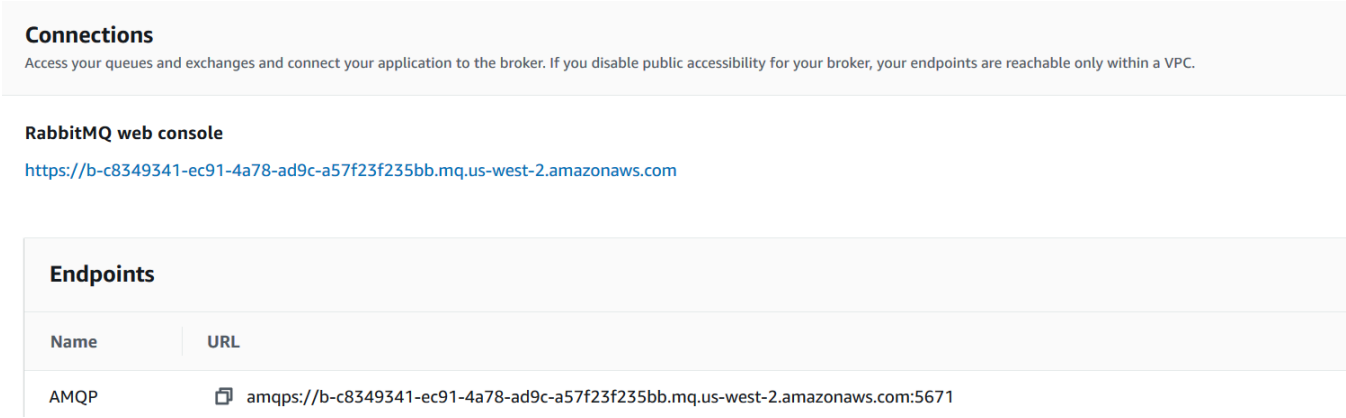

- For Amazon MQ for ActiveMQ brokers, in the **Users** section, the users associated with the broker

> **Important**
>
> Managing users via the AWS Management Console and the Amazon MQ API is not supported for Amazon MQ for RabbitMQ brokers.

Accessing the broker web console without public accessibility

If you disable public accessibility for your broker, you must perform the following steps to be able to access your broker's web console.

> **Note**
>
> The names of the VPCs and security groups are specific to the following example.

Prerequisites

To perform the following steps, you must configure the following:

- **VPCs**
  - The VPC without an internet gateway, to which the Amazon MQ broker is attached, named `private-vpc`.
  - A second VPC, with an internet gateway, named `public-vpc`.

- Both VPCs must be connected (for example, using VPC peering) so that the Amazon EC2 instances in the public VPC can communicate with the EC2 instances in the private VPC.
- If you use VPC peering, the route tables for both VPCs must be configured for the peering connection.

- **Security Groups**
  - The security group used to create the Amazon MQ broker, named `private-sg`.
  - A second security group used for the EC2 instance in the `public-vpc` VPC, named `public-sg`.
  - `private-sg` must allow inbound connections from `public-sg`. We recommend restricting this security group to port 8162 for ActiveMQ, and port 443 for RabbitMQ.
  - `public-sg` must allow inbound connections from your machine on port 22.

To Access a broker's web console of a Broker without Public Accessibility

1. Create a Linux EC2 instance in `public-vpc` (with a public IP, if necessary).
2. To verify that your VPC is configured correctly, establish an `ssh` connection to the EC2 instance and use the `curl` command with the URI of your broker.
3. From your machine, create an `ssh` tunnel to the EC2 instance using the path to your private key file and the IP address of your public EC2 instance. For example:

   ```
   ssh -i ~/.ssh/id_rsa -N -C -q -f -D 8080 ec2-user@203.0.113.0
   ```

   A forward proxy server is started on your machine.
4. Install a proxy client such as FoxyProxy on your machine.
5. Configure your proxy client using the following settings:

   - For proxy type, specify `SOCKS5`.
   - For IP address, DNS name, and server name, specify `localhost`.
   - For port, specify `8080`.
   - Remove any existing URL patterns.
   - For the URL pattern, specify `*.mq.*.amazonaws.com*`
   - For the connection type, specify `HTTP(S)`.

When you enable your proxy client, you can access the web console on your machine.

Rebooting an Amazon MQ broker

To apply a new configuration to a broker, you can reboot the broker.

> **Note**
>
> If your ActiveMQ broker becomes unresponsive, you can reboot it to recover from a faulty state.

The following example shows how you can reboot an Amazon MQ broker using the AWS Management Console.

To Reboot an Amazon MQ Broker

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**).
3. On the ***MyBroker*** page, choose **Actions**, **Reboot broker**.

> **Important**
>
> Single instance brokers will be offline while being rebooted. Cluster brokers will be available, but each node is rebooted one at a time.

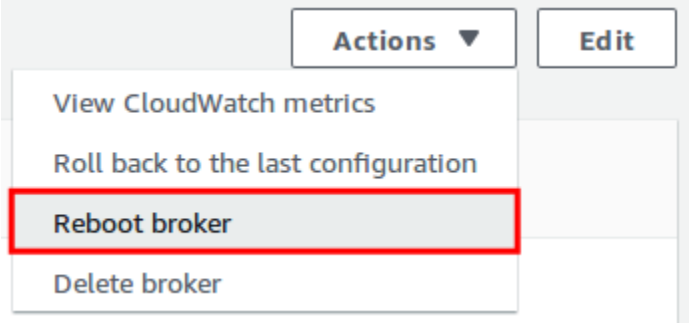

4. In the **Reboot broker** dialog box, choose **Reboot**.

Rebooting a broker takes about 5 minutes. If the reboot includes instance size changes or is performed on a broker with high queue depth, the rebooting process can take longer.

Deleting an Amazon MQ broker

If you don't use an Amazon MQ broker (and don't foresee using it in the near future), it is a best practice to delete it from Amazon MQ to reduce your AWS costs.

The following example shows how you can delete a broker using the AWS Management Console.

Deleting an Amazon MQ broker

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Delete**.
3. In the **Delete *`MyBroker`*?** dialog box, type `delete` and then choose **Delete**.

   Deleting a broker takes about 5 minutes.

Managing Amazon MQ broker configurations

A **configuration** contains all of the settings for your broker. You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers

Amazon MQ broker configuration lifecycle

Making changes to a configuration revision or an ActiveMQ user does *not* apply the changes immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.

The following diagram illustrates the configuration lifecycle.

> **Important**
>
> The next scheduled maintenance window triggers a reboot. If the broker is rebooted before the next scheduled maintenance window, the changes are applied after the reboot.

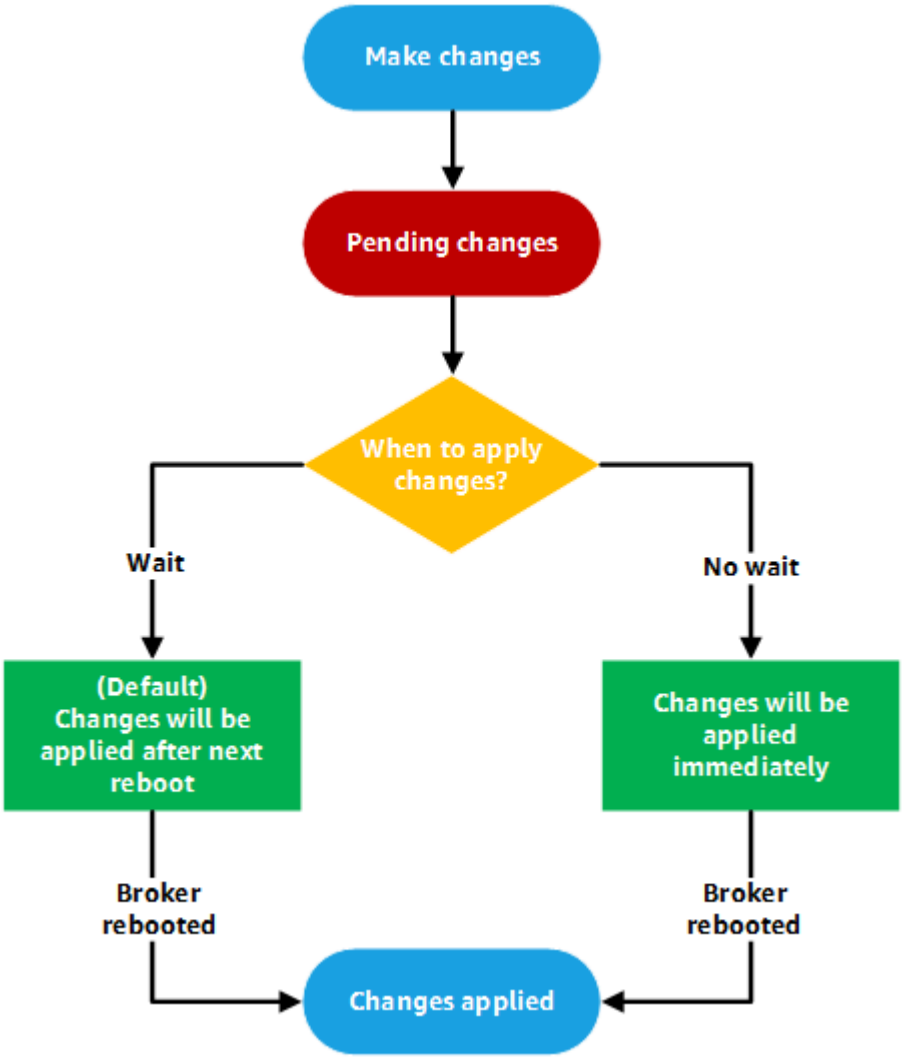

For *ActiveMQ*, a configuration contains all of the settings for your broker in XML format (similar to ActiveMQ's `activemq.xml` file). For more information on creating, applying, and editing ActiveMQ broker configurations, see Creating and applying broker configurations.

For *RabbitMQ*, a configuration constains all the settings for your broker in Cuttlefish format. For more information on creating, applying, and editing RabbitMQ broker confirgurations, see Creating and applying broker configurations

Instance types

The combined description of the broker instance *class* (`m5`, `t3`) and *size* (`large`, `micro`) is a *broker instance type* (for example, `mq.m5.large`). The following table lists the available Amazon MQ broker instance types for each supported engine type.

Topics

Amazon MQ for ActiveMQ instance types

> **Important**
>
> You can use Amazon EBS only with the `mq.m5` broker instance type family. For more information, see Storage.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.t2.micro`	1	1	Low	Use the `mq.t2.micro` instance type for basic evaluation of Amazon MQ. This instance type (single-instance brokers only) qualifies for the AWS Free Tier. **Note** Using the `mq.t2.micro` instance type is subject

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				to *CPU credits and baseline performance*—with the ability to *burst* above the baseline level (for more information, see the `CpuCreditBalance` metric). If your application requires *fixed performance*, consider using an `mq.m5.large` instance type.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.t3.micro`	2	1	Low	Use the `mq.t3.micro` instance type for basic evaluation of Amazon MQ. This instance type (single-instance brokers only) qualifies for the AWS Free Tier.
`mq.m4.large`	2	8	Moderate	Use the `mq.m4.large` instance type for compatibility with existing broker deployments. We recommend using an `mq.m5.*` instance for new brokers.
`mq.m5.large`	2	8	High	Use the `mq.m5.large` instance for regular development, testing, and production workloads.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.m5.xlarge`	4	16	High	Use the `mq.m5.xlarge` , `mq.m5.2xlarge` , and `mq.m5.4xlarge` instance types for regular development, testing and production workloads that require high throughput. **Note** When your system uses persistent messages, its throughput depends on how quickly messages are consumed.
`mq.m5.2xlarge`	8	32	High	
`mq.m5.4xlarge`	16	64	High	

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				If messages aren't consumed immediately, using larger instance types with persistent messages might not improve system throughput. In this case, we recommend setting the `concurrentStoreAndDispatchQueues` attribute to `false`. For more information, see Disable

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				Concurrent Store and Dispatch for Queues with Slow Consumers.

For more information about throughput considerations, see Choose the Correct Broker Instance Type for the Best Throughput.

Amazon MQ for RabbitMQ instance types

> **Important**
>
> You cannot downgrade a broker from an `mq.m5.` instance type to a `mq.t3.micro` instance type.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.t3.micro`	2	1	Low	Use the `mq.t3.micro` instance type for basic evaluation of Amazon MQ. This instance type (single-instance brokers

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				only) qualifies for the AWS Free Tier. **Important** The `mq.t3.micro` instance type does not support cluster deployment.
`mq.m5.large`	2	8	High	Use the `mq.m5.large` instance for regular development, testing, and production workloads.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.m5.xlarge`	4	16	High	Use the `mq.m5.xlarge` , `mq.m5.2xlarge` , and `mq.m5.4xlarge` instance types for regular development, testing and production workloads that require high throughput.
`mq.m5.2xlarge`	8	32	High	
`mq.m5.4xlarge`	16	64	High	

Tagging resources

Amazon MQ supports resource tagging to help track your cost allocation. You can tag resources when creating them, or by viewing the details of that resource.

Topics

- Tagging for Cost Allocation
- Managing Tags in the Amazon MQ Console
- Managing Using Amazon MQ API Actions

Tagging for Cost Allocation

To organize and identify your Amazon MQ resources for cost allocation, you can add metadata *tags* that identify the purpose of a broker or configuration. This is especially useful when you have many brokers. You can use cost allocation tags to organize your AWS bill to reflect your own cost structure. To do this, sign up to get your AWS account bill to include the tag keys and values. For more information, see Setting Up a Monthly Cost Allocation Report in the *AWS Billing User Guide*.

For instance, you could add tags that represent the cost center and purpose of your Amazon MQ resources:

Resource	Key	Value
Broker1	Cost Center	34567
	Stack	Production
Broker2	Cost Center	34567
	Stack	Production
Broker3	Cost Center	12345
	Stack	Development

This tagging scheme allows you to group two brokers performing related tasks in the same cost center, while tagging an unrelated broker with a different cost allocation tag.

Managing Tags in the Amazon MQ Console

Adding Tags to New Resources

Amazon MQ lets you to add tags to resources as they are created. You can quickly add tags to the resources you are creating in the Amazon MQ console.

To add tags as you create a new broker:

1. From the **Create a broker** page, select **Additional settings**.
2. Under **Tags**, select **Add tag**.
3. Enter a **Key** and **Value** pair.

4. (Optional) Select **Add tag** to add multiple tags to your broker.
5. Select **Create broker**.

To add tags as you create a configuration:

1. From the **Create configuration** page, select **Advanced**.
2. Under **Tags** on the **Create configuration** page, select **Add tag**.
3. Enter a **Key** and **Value** pair.
4. (Optional) Select **Add tag** to add multiple tags to your configuration.
5. Select **Create configuration**.

Viewing and Managing Tags for Existing Resources

Amazon MQ allows you to view and manage the tags for your resources in the Amazon MQ console. You can manage tags for an individual resource by editing the tags on the details page for that resource. To edit tags on Amazon MQ resources:

1. Select either **Brokers** or **Configurations** in the Amazon MQ console.

   Under the **Tags** section, review the existing tags for that resource.
2. To add new or manage existing tags, select **Edit** (or **Create tag** if have no existing tags).
3. Update tags for your resource:

   - To modify existing tags, edit the **Key** and **Value**.
   - To remove existing tags, select **Remove**.
   - To add a new tag, select **Add tag** and enter a **Key** and **Value**.
4. Select **Save**.

Managing Using Amazon MQ API Actions

Amazon MQ allows you to view and manage the tags of your resources using the REST API.

For more information, see the Amazon MQ REST API Reference.

Working with Amazon MQ for ActiveMQ

Amazon MQ makes it easy to create a message broker with the computing and storage resources that fit your needs. You can create, manage, and delete brokers using the AWS Management Console, Amazon MQ REST API, or the AWS Command Line Interface.

This section describes the basic elements of a message broker for ActiveMQ and RabbitMQ engine types, lists available Amazon MQ broker instance types and their statuses, and provides an overview of broker architecture and configuration options.

To learn about Amazon MQ REST APIs, see the *Amazon MQ REST API Reference*.

Topics

ActiveMQ engine

This section describes the basic elements of an ActiveMQ broker, provides an overview of ActiveMQ broker architecture options, explains broker configuration parameters, and offers a working example using Java Message Service (JMS).

Topics

Basic elements

This section introduces key concepts essential to understanding ActiveMQ on Amazon MQ.

Topics

- Broker
- Broker instance types
- Configuration
- User
- Storage

Broker

A *broker* is a message broker environment running on Amazon MQ. It is the basic building block of Amazon MQ. The combined description of the broker instance *class* (`m5`, `t3`) and *size* (`large`, `micro`) is a *broker instance type* (for example, `mq.m5.large`). For more information, see Broker instance types.

- A *single-instance broker* is comprised of one broker in one Availability Zone. The broker communicates with your application and with an Amazon EBS or Amazon EFS storage volume.
- An *active/standby broker* is comprised of two brokers in two different Availability Zones, configured in a *redundant pair*. These brokers communicate synchronously with your application, and with Amazon EFS.

For more information, see Broker Architecture.

You can enable *automatic minor version upgrades* to new minor versions of the broker engine, as Apache releases new versions. Automatic upgrades occur during the *maintenance window* defined by the day of the week, the time of day (in 24-hour format), and the time zone (UTC by default).

For information about creating and managing brokers, see the following:

- Creating and configuring a broker
- Brokers
- Broker statuses

Supported wire-level protocols

You can access your brokers by using any programming language that ActiveMQ supports and by enabling TLS explicitly for the following protocols:

- AMQP
- MQTT
- MQTT over WebSocket
- OpenWire
- STOMP
- STOMP over WebSocket

Attributes

An ActiveMQ broker has several attributes, for example:

- A name (`MyBroker`)
- An ID (`b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`)
- An Amazon Resource Name (ARN) (`arn:aws:mq:us-east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`)
- An ActiveMQ Web Console URL (`https://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:8162`)

  For more information, see Web Console in the Apache ActiveMQ documentation.

  > **Important**
  >
  > If you specify an authorization map which doesn't include the `activemq-webconsole` group, you can't use the ActiveMQ Web Console because the group isn't authorized to send messages to, or receive messages from, the Amazon MQ broker.

- Wire-level protocol endpoints:
  - `amqp+ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:5671`
  - `mqtt+ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:8883`
  - `ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61617`

> **Note**
>
> This is an OpenWire endpoint.

- `stomp+ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61614`
- `wss://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61619`

For more information, see Configuring Transports in the Apache ActiveMQ documentation.

> **Note**
>
> For an active/standby broker, Amazon MQ provides two ActiveMQ Web Console URLs, but only one URL is active at a time. Likewise, Amazon MQ provides two endpoints for each wire-level protocol, but only one endpoint is active in each pair at a time. The `-1` and `-2` suffixes denote a redundant pair.

For a full list of broker attributes, see the following in the *Amazon MQ REST API Reference*:

- REST Operation ID: Broker
- REST Operation ID: Brokers
- REST Operation ID: Broker Reboot

Broker instance types

> **Important**
>
> You can use Amazon EBS only with the `mq.m5` broker instance type family. For more information, see Storage.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.t2.micro`	1	1	Low	Use the `mq.t2.micro` instance type for basic evaluation of Amazon MQ. This instance type (single-instance brokers only) qualifies for the AWS Free Tier. **Note** Using the `mq.t2.micro` instance type is subject to *CPU credits and baseline performance*—with the ability to *burst* above the baseline level (for

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				more information, see the CpuCreditBalance metric). If your application requires *fixed performance*, consider using an `mq.m5.large` instance type.
`mq.t3.micro`	2	1	Low	Use the `mq.t3.micro` instance type for basic evaluation of Amazon MQ. This instance type (single-instance brokers only) qualifies for the AWS Free Tier.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.m4.large`	2	8	Moderate	Use the `mq.m4.large` instance type for compatibility with existing broker deployments. We recommend using an `mq.m5.*` instance for new brokers.
`mq.m5.large`	2	8	High	Use the `mq.m5.large` instance for regular development, testing, and production workloads.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.m5.xlarge`	4	16	High	Use the `mq.m5.xlarge` , `mq.m5.2xlarge` , and `mq.m5.4xlarge` instance types for regular development, testing and production workloads that require high throughput. **Note** When your system uses persistent messages, its throughput depends on how quickly messages are consumed.
`mq.m5.2xlarge`	8	32	High	
`mq.m5.4xlarge`	16	64	High	

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				If messages aren't consumed immediately, using larger instance types with persistent messages might not improve system throughput. In this case, we recommend setting the `concurrentStoreAndDispatchQueues` attribute to `false`. For more information, see Disable

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				Concurrent Store and Dispatch for Queues with Slow Consumers.

For more information about throughput considerations, see Choose the Correct Broker Instance Type for the Best Throughput.

Configuration

A *configuration* contains all of the settings for your ActiveMQ broker, in XML format (similar to ActiveMQ's `activemq.xml` file). You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers.

> **Important**
>
> Making changes to a configuration does *not* apply the changes to the broker immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle. Currently, you can't delete a configuration.

For information about creating, editing, and managing configurations, see the following:

- Creating and applying broker configurations
- Configurations
- Amazon MQ Broker Configuration Parameters

To keep track of the changes you make to your configuration, you can create *configuration revisions*. For more information, see Creating and applying broker configurations.

Attributes

A broker configuration has several attributes, for example:

- A name (`MyConfiguration`)
- An ID (`c-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`)
- An Amazon Resource Name (ARN) (`arn:aws:mq:us-east-2:123456789012:configuration:c-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`)

For a full list of configuration attributes, see the following in the *Amazon MQ REST API Reference*:

- REST Operation ID: Configuration
- REST Operation ID: Configurations

For a full list of configuration revision attributes, see the following:

- REST Operation ID: Configuration Revision
- REST Operation ID: Configuration Revisions

User

An ActiveMQ *user* is a person or an application that can access the queues and topics of an ActiveMQ broker. You can configure users to have specific permissions. For example, you can allow some users to access the ActiveMQ Web Console.

A *group* is a semantic label. You can assign a group to a user and configure permissions for groups to send to, receive from, and administer specific queues and topics.

> **Important**
>
> Making changes to a user does *not* apply the changes to the user immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.

For information about users and groups, see the following in the Apache ActiveMQ documentation:

- Authorization
- Authorization Example

For information about creating, editing, and deleting ActiveMQ users, see the following:

- Creating and managing ActiveMQ broker users
- Users

Attributes

For a full list of user attributes, see the following in the *Amazon MQ REST API Reference*:

- REST Operation ID: User
- REST Operation ID: Users

Storage

Amazon MQ for ActiveMQ supports Amazon Elastic File System (EFS) and Amazon Elastic Block Store (EBS). By default, ActiveMQ brokers use Amazon EFS for broker storage. To take advantage of high durability and replication across multiple Availability Zones, use Amazon EFS. To take advantage of low latency and high throughput, use Amazon EBS.

> **Important**
>
> - You can use Amazon EBS only with the `mq.m5` broker instance type family.
> - Although you can change the *broker instance type*, you can't change the *broker storage type* after you create the broker.
> - Amazon EBS replicates data within a single Availability Zone and doesn't support the ActiveMQ active/standby deployment mode.

Differences between Storage Types

The following table provides a brief overview of the differences between in-memory, Amazon EFS, and Amazon EBS storage types for ActiveMQ brokers.

Storage Type	Persistence	Example Use Case	Approximate Maximum Number of Messages Enqueued per Producer, per Second (1KB Message)	Replication
In-memory	Non-persistent	• Stock quotes • Location data updates • Frequently changed data	5,000	None
Amazon EBS	Persistent	• High volumes of text • Order processing	500	Multiple copies within a single Availability Zone (AZ)
Amazon EFS	Persistent	Financial transactions	80	Multiple copies across multiple AZs

In-memory message storage provides the lowest latency and the highest throughput. However, messages are lost during instance replacement or broker restart.

Amazon EFS is designed to be highly durable, replicated across multiple AZs to prevent the loss of data resulting from the failure of any single component or an issue that affects the availability of an AZ. Amazon EBS is optimized for throughput and replicated across multiple servers within a single AZ.

Broker architecture

Amazon MQ for ActiveMQ brokers can be created as *single-instance brokers* or *active/standby brokers*. For both deployment modes, Amazon MQ provides high durability by storing its data redundantly.

> **Note**
>
> Amazon MQ uses Apache KahaDB as its data store. Other data stores, such as JDBC and LevelDB, aren't supported.

You can access your brokers by using any programming language that ActiveMQ supports and by enabling TLS explicitly for the following protocols:

- AMQP
- MQTT
- MQTT over WebSocket
- OpenWire
- STOMP
- STOMP over WebSocket

Topics

Amazon MQ single-instance broker

A *single-instance broker* is comprised of one broker in one Availability Zone. The broker communicates with your application and with an Amazon EBS or Amazon EFS storage volume. Amazon EFS storage volumes are designed to provide the highest level of durability and availability by storing data redundantly across multiple Availability Zones (AZs). Amazon EBS provides block level storage optimized for low-latency and high throughput. For more information about storage options, see Storage.

The following diagram illustrates a single-instance broker with Amazon EFS storage replicated across multiple AZs.

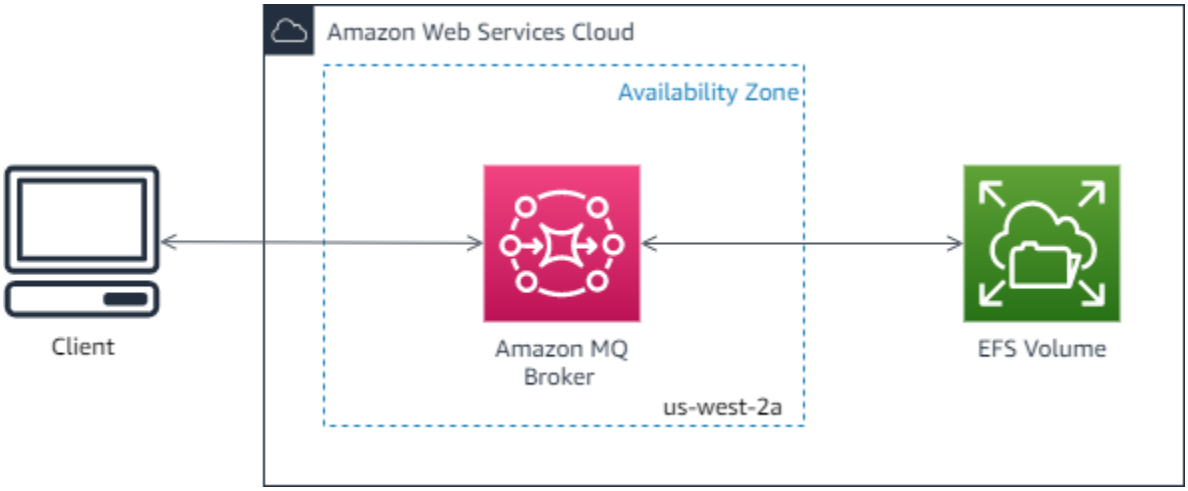

The following diagram illustrates a single-instance broker with Amazon EBS storage replicated across multiple servers within a single AZ.

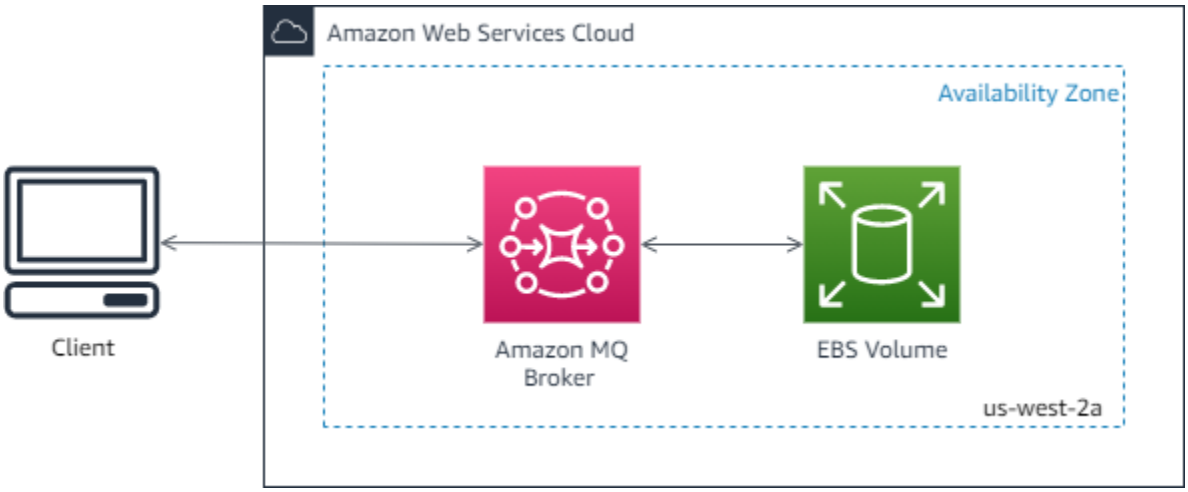

Amazon MQ active/standby broker for high availability

An *active/standby broker* is comprised of two brokers in two different Availability Zones, configured in a *redundant pair*. These brokers communicate synchronously with your application, and with Amazon EFS. Amazon EFS storage volumes are designed to provide the highest level of durability, and availability by storing data redundantly across multiple Availability Zones (AZs). For more information, see Storage.

Usually, only one of the broker instances is active at any time, while the other broker instance is on standby. If one of the broker instances malfunctions or undergoes maintenance, it takes Amazon MQ a short while to take the inactive instance out of service. This allows the healthy standby instance to become active and to begin accepting incoming communications. When you reboot a broker, the failover takes only a few seconds.

For an active/standby broker, Amazon MQ provides two ActiveMQ Web Console URLs, but only one URL is active at a time. Likewise, Amazon MQ provides two endpoints for each wire-level protocol, but only one endpoint is active in each pair at a time. The `-1` and `-2` suffixes denote a redundant pair. For wire-level protocol endpoints, you can allow your application to connect to either endpoint by using the Failover Transport.

The following diagram illustrates an active/standby broker with Amazon EFS storage replicated across multiple AZs.

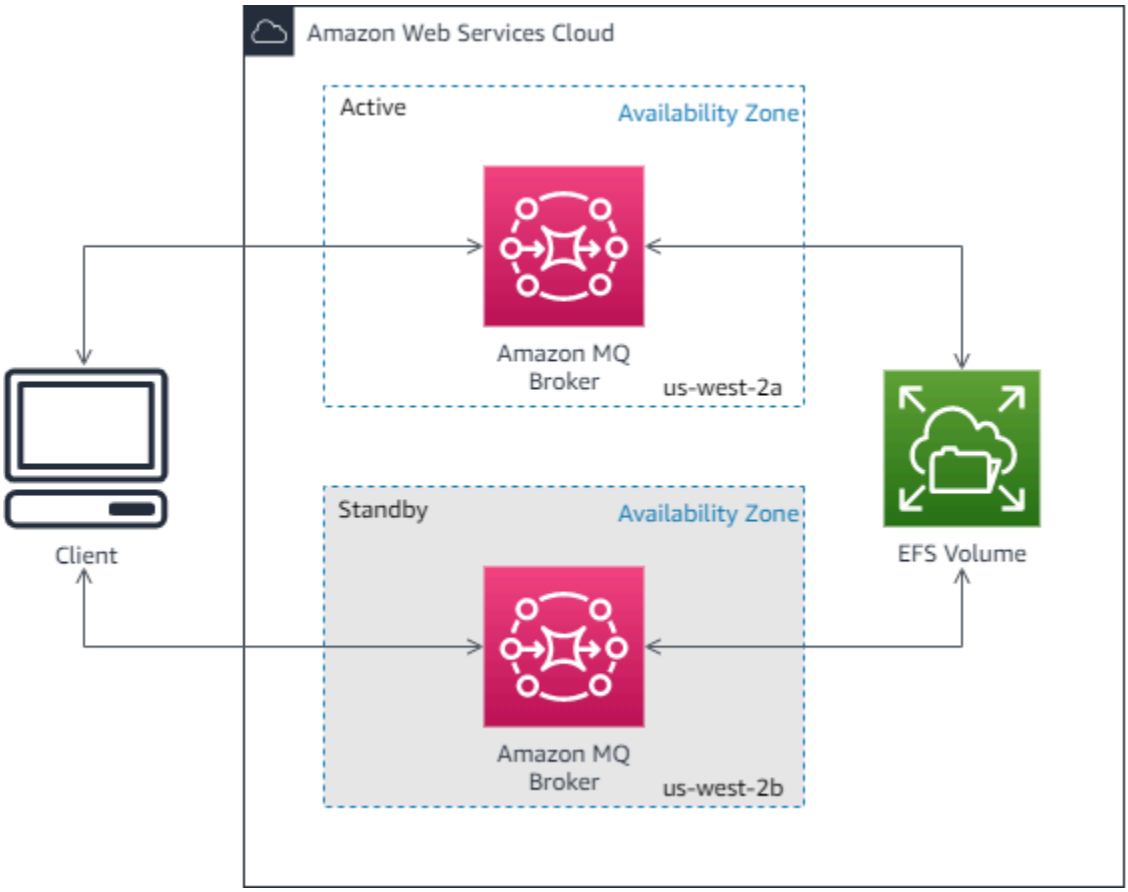

Amazon MQ Network of brokers

Amazon MQ supports ActiveMQ's network of brokers feature.

A *network of brokers* is comprised of multiple simultaneously active single-instance brokers or active/standby brokers. You can configure networks of brokers in a variety of topologies (for example, *concentrator*, *hub-and-spokes*, *tree*, or *mesh*), depending on your application's needs, such as high availability and scalability. For instance, a hub and spoke network of brokers can increase resiliency, preserving messages if one broker is not reachable. A network of brokers with a concentrator topology can collect messages from a larger number of brokers accepting incoming

messages, and concentrate them to more central brokers, to better handle the load of many incoming messages.

For a tutorial and detailed configuration information, see the following:

- Creating and Configuring a Network of Brokers
- Configure Your Network of Brokers Correctly
- `networkConnector`
- `networkConnectionStartAsync`
- Networks of Brokers in the ActiveMQ documentation

The following are benefits of using a network of brokers:

- Creating a network of brokers allows you to increase your aggregate throughput and maximum producer and consumer connection count by adding broker instances.
- You can ensure better availability by allowing your producers and consumers to be aware of multiple active broker instances. This allows them to reconnect to a new instance if the one they're currently connected to becomes unavailable.
- Because producers and consumers can reconnect to another node in the network of brokers immediately, and because there's no need to wait for a standby broker instance to become promoted, client reconnection within a network of brokers is faster than for an active/standby broker for high availability.

Topics

- How does a network of brokers work?
- How Does a Network of Brokers Handle Credentials?
- Sample blueprints
- Network of brokers topologies
- Cross region
- Dynamic Failover With Transport Connectors

How does a network of brokers work?

Amazon MQ supports the ActiveMQ network of brokers feature in a number of ways. First, you can edit the parameters within each broker's configuration to create a network of brokers, just

as you would with native ActiveMQ. Second, Amazon MQ has sample blueprints that use AWS CloudFormation to automate the creation of a network of brokers. You can deploy these sample blueprints directly from the Amazon MQ console, or you can edit the related AWS CloudFormation templates to create your own topologies and configurations.

A network of brokers is established by connecting one broker to another using network connectors. Once connected, these brokers provide message forwarding. For instance, if *Broker1* establishes a network connector to *Broker2*, messages on *Broker1* are forwarded to *Broker2* if there is a consumer on that broker for the queue or topic. If the network connector is configured as `duplex`, messages are also forwarded from *Broker2* to *Broker1*. Network connectors are configured in the broker *configuration*. See, Configuration. For instance, here is an example `networkConnector` entry in a broker configuration:

```
<networkConnectors>
  <networkConnector name="connector_1_to_2" userName="myCommonUser" duplex="true"
    uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

A network of brokers ensures that messages flow from one broker instance to another, forwarding messages only to the broker instances that have corresponding consumers. For the benefit of broker instances adjacent to each other within the network, ActiveMQ sends messages to *advisory topics* about producers and consumers connecting to and disconnecting from the network. When a broker instance receives information about a consumer that consumes from a particular destination, the broker instance begins to forward messages. For more information, see Advisory Topics in the ActiveMQ documentation.

How Does a Network of Brokers Handle Credentials?

For broker A to connect to broker B in a network, broker A must use valid credentials, like any other producer or consumer. Instead of providing a password in broker A's `<networkConnector>` configuration, you must first create a user on broker A with the same values as another user on broker B (these are *separate, unique* users that share the same username and password values). When you specify the `userName` attribute in the `<networkConnector>` configuration, Amazon MQ will add the password automatically at runtime.

> **Important**
>
> Don't specify the `password` attribute for the `<networkConnector>`. We don't recommend storing plaintext passwords in broker configuration files, because this makes the passwords visible in the Amazon MQ console. For more information, see Configure Network Connectors for Your Broker.

Brokers must be in the same VPC or in peered VPCs. For more information, see Prerequisites in the Creating and Configuring a Network of Brokers tutorial.

Sample blueprints

To get started using a Network of Brokers, Amazon MQ provides sample blueprints. These sample blueprints create a Network of Brokers deployment, and all related resources using, AWS CloudFormation. The two sample blueprints available are:

1. Mesh network of single instance brokers
2. Mesh network of active/standby brokers

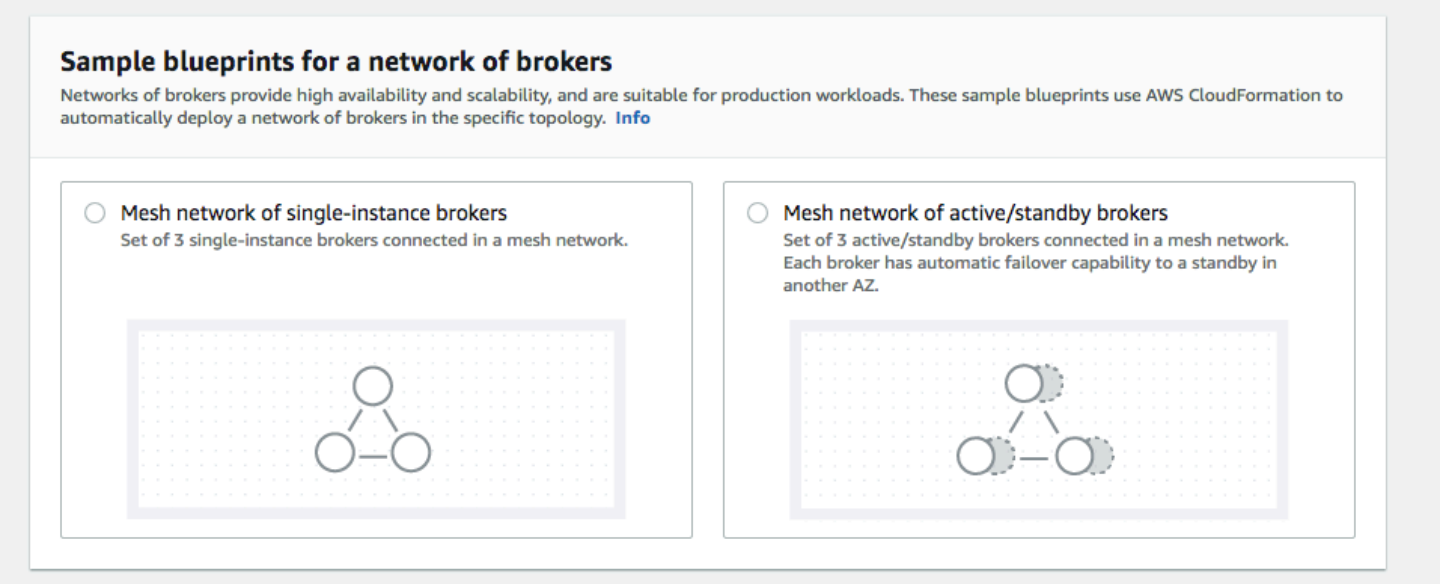

From the **Create brokers** page, select one of the sample blueprints and choose **Next**. Once the resources have been created, review the generated brokers and their configurations in the Amazon MQ console.

By creating brokers and configuring different `networkConnector` elements in the broker configurations, you can create a network of brokers in many different topologies. For more information on configuring a network of brokers, see Networks of Brokers in the ActiveMQ documentation.

Network of brokers topologies

By deploying brokers, and then configuring `networkConnector` entries in their configurations, you can build a network of brokers using different network topologies. A network connector provides on-demand message forwarding between connected brokers. Connections can be configured as duplex, where messages are forwarded both ways between brokers, or not duplex, where the forwarding only propagates from one broker to the other. For example, if we have a duplex connection between *Broker1* and *Broker2*, messages will be forwarded from each to the other if there is a consumer.

With a duplex network connector, messages are forwarded from each broker to the other. These are forwarded *on-demand*: if there is a consumer on *Broker2* for a message on *Broker1*, the message is forwarded. Similarly, if there is a consumer on *Broker1* for a message on *Broker2* the message is also forwarded.

For non-duplex connections, messages are forwarded only from one broker to the other. In this example, if there is a consumer on *Broker2* for a message on *Broker1*, the message is forwarded. But messages will not be forwarded from *Broker2* to *Broker1*.

Using both duplex and non-duplex network connectors, it is possible to build a network of brokers in any number of network topologies.

> **Note**
>
> In each of the network topology examples, the `networkConnector` elements reference the endpoint of the brokers they connect to. Replace the broker endpoint entries in the `uri` attributes with the endpoints of your brokers. See, Listing brokers and viewing broker details.

Mesh topology

A mesh topology provides multiple brokers that are all connected to each other. This simple example connects three single-instance brokers, but you can configure more brokers as a mesh.

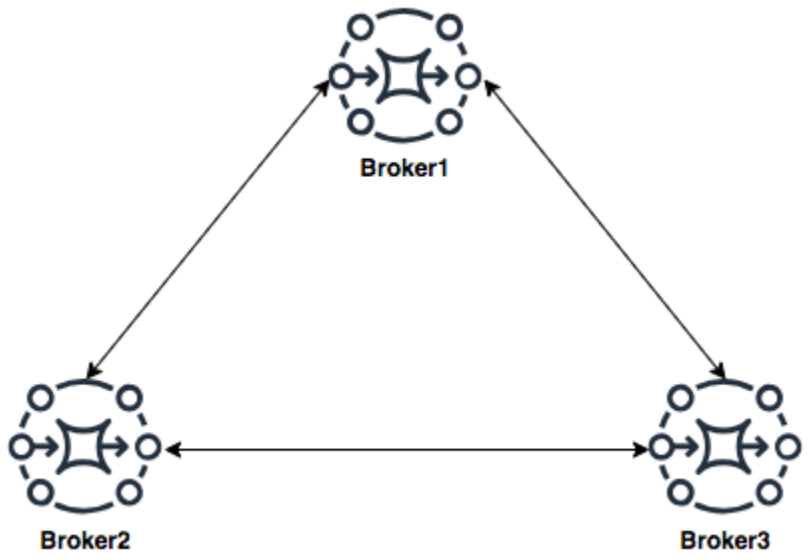

This topology, and one that includes a mesh of *active/standby* pairs of brokers, can be created using sample blueprints in the Amazon MQ console. You can create these sample blueprint deployment to see a working network of brokers, and review how they are configured.

You can configure a three broker mesh network like this by adding a network connector to *Broker1* that makes duplex connections to both *Broker2* and *Broker3*, and a single duplex connection between *Broker2* and *Broker3*.

Network connectors for Broker1:

```
<networkConnectors>
```

```
    <networkConnector name="connector_1_to_2" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="connector_1_to_3" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-743c885d-2244-4c95-af67-a85017ff234e-3.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Network connectors for Broker2:

```
<networkConnectors>
    <networkConnector name="connector_2_to_3" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-743c885d-2244-4c95-af67-a85017ff234e-3.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

By adding the above connectors to the configurations of *Broker1* and *Broker2*, you can create a mesh between these three brokers that forwards message between all the brokers on demand. For more information, see Amazon MQ Broker Configuration Parameters.

Hub and Spoke Topology

In a hub and spoke topology, messages are preserved if there is a disruption to any broker on a spoke. Messages are forwarded throughout, and only the central *Broker1* is critical to the network's operation.

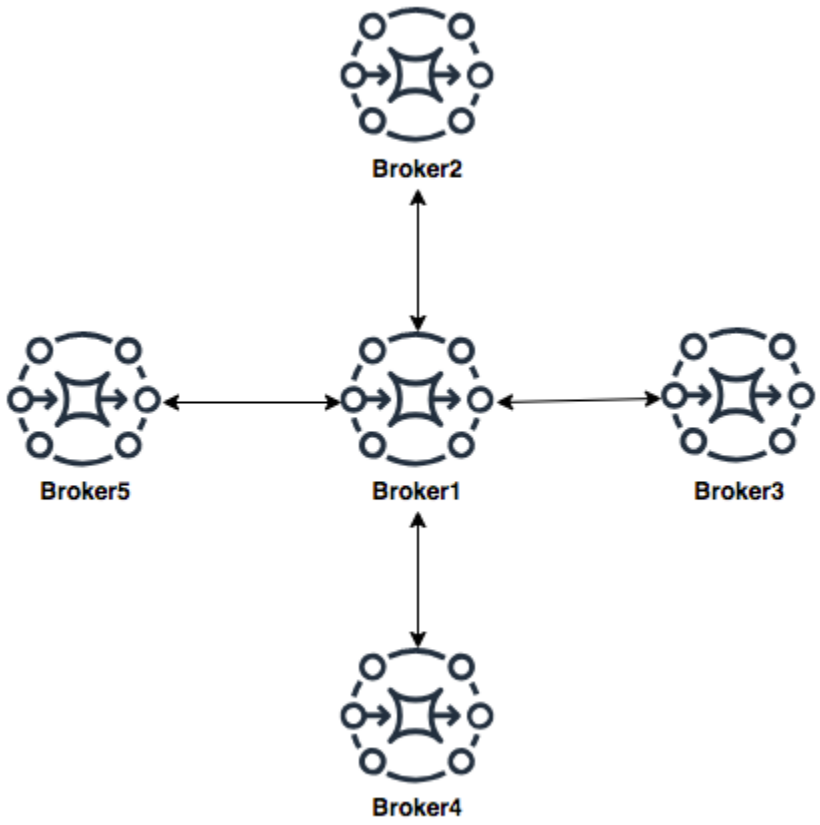

To configure the hub and spoke network of brokers in this example, you could add a `networkConnector` to each of the brokers on the spokes in the configuration of *Broker1*.

```
<networkConnectors>
    <networkConnector name="connector_hub_and_spoke_2" userName="myCommonUser"
 duplex="true"
        uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="connector_hub_and_spoke_3" userName="myCommonUser"
 duplex="true"
        uri="static:(ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="connector_hub_and_spoke_4" userName="myCommonUser"
 duplex="true"
        uri="static:(ssl://b-743c885d-2244-4c95-af67-a85017ff234e-3.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="connector_hub_and_spoke_5" userName="myCommonUser"
 duplex="true"
```

```
        uri="static:(ssl://b-62a7fb31-d51c-466a-a873-905cd660b553-4.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Concentrator topology

In this example topology, the three brokers on the bottom can handle a large number of connections, and those messages are concentrated to *Broker1* and *Broker2*. Each of the other brokers has a non-duplex connection to the more central brokers. To scale the capacity of this topology, you can add additional brokers that receive messages and concentrate those messages in *Broker1* and *Broker2*.

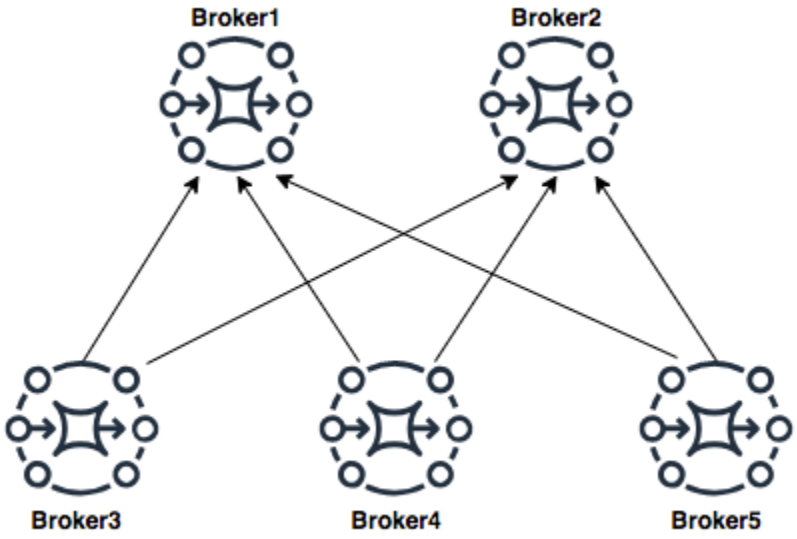

To configure this topology, each of the brokers on the bottom would contain a network connector to each of the brokers they are concentrating messages to.

Network connectors for Broker3:

```
<networkConnectors>
    <networkConnector name="3_to_1" userName="myCommonUser" duplex="false"
        uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="3_to_2" userName="myCommonUser" duplex="false"
        uri="static:(ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Network connectors for Broker4:

```
<networkConnectors>
    <networkConnector name="4_to_1" userName="myCommonUser" duplex="false"
        uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="4_to_2" userName="myCommonUser" duplex="false"
        uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Network connectors for Broker5:

```
<networkConnectors>
    <networkConnector name="5_to_1" userName="myCommonUser" duplex="false"
        uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="5_to_2" userName="myCommonUser" duplex="false"
        uri="static:(ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Cross region

To configure a network of brokers that spans AWS regions, deploy brokers in those regions, and configure network connectors to the endpoints of those brokers.

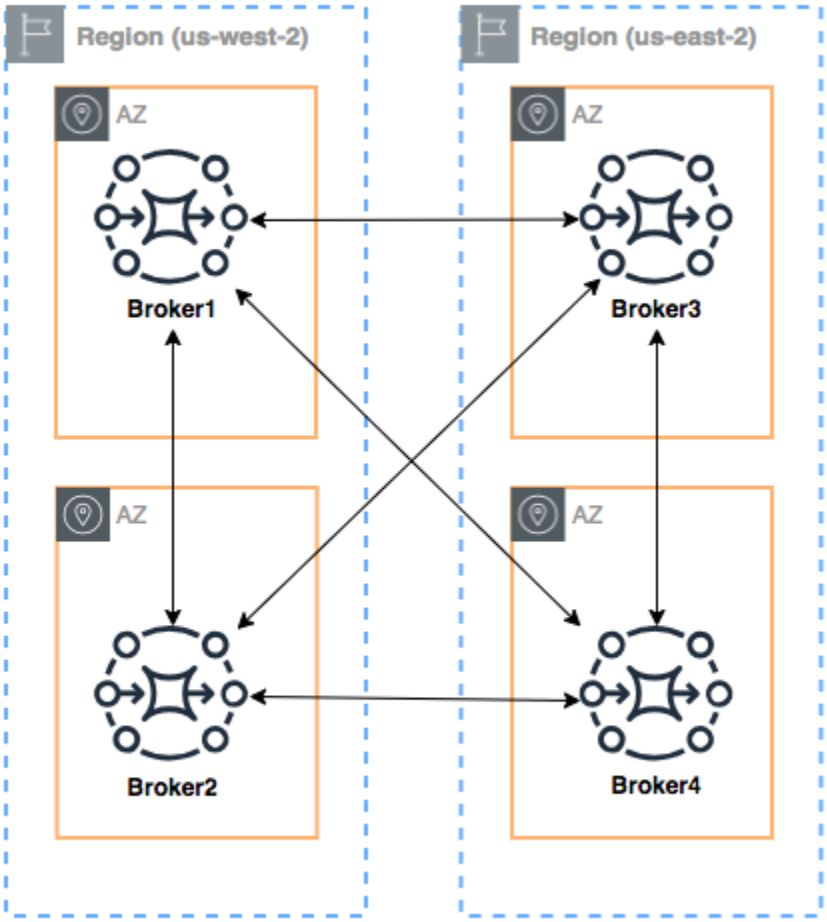

To configure a network of brokers like this example, you could add `networkConnectors` entries to the configurations of *Broker1* and *Broker4* that reference the wire-level endpoints of those brokers.

Network connectors for Broker1:

```
<networkConnectors>
    <networkConnector name="1_to_2" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="1_to_3" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-743c885d-2244-4c95-af67-a85017ff234e-3.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="1_to_4" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-62a7fb31-d51c-466a-a873-905cd660b553-4.mq.us-
east-2.amazonaws.com:61617)"/>
```

```
</networkConnectors>
```

Network connector for Broker2:

```
<networkConnectors>
    <networkConnector name="2_to_3" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-743c885d-2244-4c95-af67-a85017ff234e-3.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Network connectors for Broker4:

```
<networkConnectors>
    <networkConnector name="4_to_3" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-743c885d-2244-4c95-af67-a85017ff234e-3.mq.us-
east-2.amazonaws.com:61617)"/>
    <networkConnector name="4_to_2" userName="myCommonUser" duplex="true"
        uri="static:(ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

Dynamic Failover With Transport Connectors

In addition to configuring `networkConnector` elements, you can configure your broker `transportConnector` options to enable dynamic failover, and to rebalance connections when brokers are added or removed from the network.

```
<transportConnectors>
  <transportConnector name="openwire" updateClusterClients="true"
 rebalanceClusterClients="true" updateClusterClientsOnRemove="true"/>
</transportConnectors>
```

In this example both `updateClusterClients` and `rebalanceClusterClients` are set to `true`. In this case clients will be provided a list of brokers in the network, and will request them to rebalance if a new broker joins.

Available options:

- `updateClusterClients`: Passes information to clients about changes in the network of broker topology.

- `rebalanceClusterClients`: Causes clients to re-balance across brokers when a new broker is added to a network of brokers.
- `updateClusterClientsOnRemove`: Updates clients with topology information when a broker leaves a network of brokers.

When `updateClusterClients` is set to true, clients can be configured to connect to a single broker in a network of brokers.

```
failover:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)
```

When a new broker connects, it will receive a list of URIs of all brokers in the network. If the connection to the broker fails, it can dynamically switch to one of the brokers provided when it connected.

For more information on failover, see Broker-side Options for Failover in the Active MQ documentation.

Amazon MQ for ActiveMQ broker configurations

A configuration contains all of the settings for your ActiveMQ broker in XML format (similar to ActiveMQ's `activemq.xml` file). You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers.

Topics

Working with Spring XML configuration files

ActiveMQ brokers are configured using Spring XML files. You can configure many aspects of your ActiveMQ broker, such as predefined destinations, destination policies, authorization policies, and

plugins. Amazon MQ controls some of these configuration elements, such as network transports and storage. Other configuration options, such as creating networks of brokers, aren't currently supported.

The full set of supported configuration options is specified in the Amazon MQ XML schemas. Download zip files of the supported schemas using the following links.

- `amazon-mq-active-mq-5.17.6.xsd.zip`
- `amazon-mq-active-mq-5.16.7.xsd.zip`
- `amazon-mq-active-mq-5.15.16.xsd.zip`

You can use these schemas to validate and sanitize your configuration files. Amazon MQ also lets you provide configurations by uploading XML files. When you upload an XML file, Amazon MQ automatically sanitizes and removes invalid and prohibited configuration parameters according to the schema.

> **Note**
>
> You can use only static values for attributes. Amazon MQ sanitizes elements and attributes that contain Spring expressions, variables, and element references from your configuration.

Creating, editing, and applying ActiveMQ broker configurations

A *configuration* contains all of the settings for your ActiveMQ broker, in XML format (similar to ActiveMQ's `activemq.xml` file). You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers. You can apply a configuration immediately or during a *maintenance window*.

For more information, see the following:

- Configuration
- Amazon MQ broker configuration lifecycle
- Amazon MQ Broker Configuration Parameters

The following example shows how you can create and apply an Amazon MQ broker configuration using the AWS Management Console.

Topics

- Create a New Configuration
- Create a New Configuration Revision
- Apply a Configuration Revision to Your Broker
- Edit a Configuration Revision

Create a New Configuration

1. Sign in to the Amazon MQ console.
2. On the left, expand the navigation panel and choose **Configurations**.

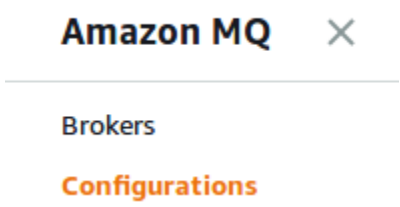

3. On the **Configurations** page, choose **Create configuration**.
4. On the **Create configuration** page, in the **Details** section, type the **Configuration name** (for example, `MyConfiguration`) and select a **Broker engine** version.

> **Note**
>
> To learn more about ActiveMQ engine versions supported by Amazon MQ for ActiveMQ, see the section called "Version management".

5. Choose **Create configuration**.

Create a New Configuration Revision

1. From the configuration list, choose ***`MyConfiguration`***.

> **Note**
>
> The first configuration revision is always created for you when Amazon MQ creates the configuration.

On the ***MyConfiguration*** page, the broker engine type and version that your new configuration revision uses (for example, **Apache ActiveMQ 5.15.16**) are displayed.

2. On the **Configuration details** tab, the configuration revision number, description, and broker configuration in XML format are displayed.

> **Note**
>
> Editing the current configuration creates a new configuration revision.

Revision 1 Auto-generated default for MyBroker-configuration on ActiveMQ 5.15.0 **Latest**

Amazon MQ configurations support a limited subset of ActiveMQ properties. Info

```
1 <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
2 <broker xmlns="http://activemq.apache.org/schema/core">
3   <!--
4   A configuration contains all of the settings for your ActiveMQ broker, in XML format
(similar to ActiveMQ's activemq.xml file).
5   You can create a configuration before creating any brokers. You can then apply the
configuration to one or more brokers.
```

3. Choose **Edit configuration** and make changes to the XML configuration.
4. Choose **Save**.

   The **Save revision** dialog box is displayed.
5. (Optional) Type `A description of the changes in this revision`.
6. Choose **Save**.

   The new revision of the configuration is saved.

> **Important**
>
> The Amazon MQ console automatically sanitizes invalid and prohibited configuration parameters according to a schema. For more information and a full list of permitted XML parameters, see Amazon MQ Broker Configuration Parameters.
> Making changes to a configuration does *not* apply the changes to the broker immediately. To apply your changes, you must wait for the next maintenance window

or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.
Currently, you can't delete a configuration.

Apply a Configuration Revision to Your Broker

1. On the left, expand the navigation panel and choose **Brokers**.

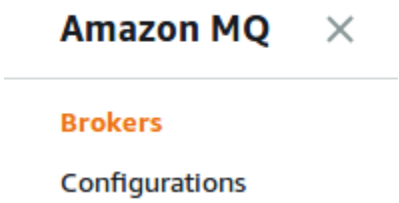

2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Edit**.
3. On the **Edit *MyBroker*** page, in the **Configuration** section, select a **Configuration** and a **Revision** and then choose **Schedule Modifications**.
4. In the **Schedule broker modifications** section, choose whether to apply modifications **During the next scheduled maintenance window** or **Immediately**.

> **Important**
> Your broker will be offline while it is being rebooted.

5. Choose **Apply**.

   Your configuration revision is applied to your broker at the specified time.

Edit a Configuration Revision

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Edit**.
3. On the ***MyBroker*** page, choose **Edit**.
4. On the **Edit *MyBroker*** page, in the **Configuration** section, select a **Configuration** and a **Revision** and then choose **Edit**.

> **Note**
>
> Unless you select a configuration when you create a broker, the first configuration revision is always created for you when Amazon MQ creates the broker.

On the ***MyBroker*** page, the broker engine type and version that the configuration uses (for example, **Apache ActiveMQ 5.15.8**) are displayed.

5. On the **Configuration details** tab, the configuration revision number, description, and broker configuration in XML format are displayed.

> **Note**
>
> Editing the current configuration creates a new configuration revision.

Revision 1 Auto-generated default for MyBroker-configuration on ActiveMQ 5.15.0 **Latest**

Amazon MQ configurations support a limited subset of ActiveMQ properties. Info

```
1 <?xml version="1.0" encoding="UTF-8" standalone="yes"?>
2 <broker xmlns="http://activemq.apache.org/schema/core">
3   <!--
4   A configuration contains all of the settings for your ActiveMQ broker, in XML format
  (similar to ActiveMQ's activemq.xml file).
5   You can create a configuration before creating any brokers. You can then apply the
  configuration to one or more brokers.
```

6. Choose **Edit configuration** and make changes to the XML configuration.
7. Choose **Save**.

   The **Save revision** dialog box is displayed.
8. (Optional) Type `A description of the changes in this revision`.
9. Choose **Save**.

   The new revision of the configuration is saved.

> **Important**
>
> The Amazon MQ console automatically sanitizes invalid and prohibited configuration parameters according to a schema. For more information and a full list of permitted XML parameters, see Amazon MQ Broker Configuration Parameters.
> Making changes to a configuration does *not* apply the changes to the broker immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.
> Currently, you can't delete a configuration.

Elements Permitted in Amazon MQ Configurations

The following is a detailed listing of the elements permitted in Amazon MQ configurations. For more information, see XML Configuration in the Apache ActiveMQ documentation.

Element
`abortSlowAckConsumerStrategy` (attributes)
`abortSlowConsumerStrategy` (attributes)
`authorizationEntry` (attributes)
`authorizationMap` (child collection elements)
`authorizationPlugin` (child collection elements)
`broker` (attributes | child collection elements)
`cachedMessageGroupMapFactory` (attributes)
`compositeQueue` (attributes | child collection elements)
`compositeTopic` (attributes | child collection elements)
`constantPendingMessageLimitStrategy` (attributes)

Element
`discarding` (attributes)
`discardingDLQBrokerPlugin` (attributes)
`fileCursor`
`fileDurableSubscriberCursor`
`fileQueueCursor`
`filteredDestination` (attributes)
`fixedCountSubscriptionRecoveryPolicy` (attributes)
`fixedSizedSubscriptionRecoveryPolicy` (attributes)
`forcePersistencyModeBrokerPlugin` (attributes)
`individualDeadLetterStrategy` (attributes)
`lastImageSubscriptionRecoveryPolicy`
`messageGroupHashBucketFactory` (attributes)
`mirroredQueue` (attributes)
`noSubscriptionRecoveryPolicy`
`oldestMessageEvictionStrategy` (attributes)
`oldestMessageWithLowestPriorityEvictionStrategy` (attributes)
`policyEntry` (attributes | child collection elements)
`policyMap` (child collection elements)
`prefetchRatePendingMessageLimitStrategy` (attributes)
`priorityDispatchPolicy`

Element	
priorityNetworkDispatchPolicy	
queryBasedSubscriptionRecoveryPolicy	(attributes)
queue (attributes)	
redeliveryPlugin	(attributes | child collection elements)
redeliveryPolicy	(attributes)
redeliveryPolicyMap	(child collection elements)
retainedMessageSubscriptionRecoveryPolicy	(child collection elements)
roundRobinDispatchPolicy	
sharedDeadLetterStrategy	(attributes | child collection elements)
simpleDispatchPolicy	
simpleMessageGroupMapFactory	
statisticsBrokerPlugin	
storeCursor	
storeDurableSubscriberCursor	(attributes)
strictOrderDispatchPolicy	
tempDestinationAuthorizationEntry	(attributes)
tempQueue	(attributes)
tempTopic	(attributes)
timedSubscriptionRecoveryPolicy	(attributes)
timeStampingBrokerPlugin	(attributes)

Element
`topic` (attributes)
`transportConnector` (attributes)
`uniquePropertyMessageEvictionStrategy` (attributes)
`virtualDestinationInterceptor` (child collection elements)
`virtualTopic` (attributes)
`vmCursor`
`vmDurableCursor`
`vmQueueCursor`

Elements and Their Attributes Permitted in Amazon MQ Configurations

The following is a detailed listing of the elements and their attributes permitted in Amazon MQ configurations. For more information, see XML Configuration in the Apache ActiveMQ documentation.

Element	Attribute
`abortSlowAckConsumerStrategy`	`abortConnection`
	`checkPeriod`
	`ignoreIdleConsumers`
	`ignoreNetworkConsumers`
	`maxSlowCount`
	`maxSlowDuration`
	`maxTimeSinceLastAck`

Element	Attribute
	name
abortSlowConsumerStrategy	abortConnection
	checkPeriod
	ignoreNetworkConsumers
	maxSlowCount
	maxSlowDuration
	name
authorizationEntry	admin
	queue
	read
	tempQueue
	tempTopic
	topic
	write
broker	advisorySupport
	allowTempAutoCreationOnSend
	cacheTempDestinations
	consumerSystemUsagePortion
	dedicatedTaskRunner
	deleteAllMessagesOnStartup

Element	Attribute
	keepDurableSubsActive
	enableMessageExpirationOnActiveDurableSubs
	maxPurgedDestinationsPerSweep
	maxSchedulerRepeatAllowed
	monitorConnectionSplits
	networkConnectorStartAsync
	offlineDurableSubscriberTaskSchedule
	offlineDurableSubscriberTimeout
	persistenceThreadPriority
	persistent
	populateJMSXUserID
	producerSystemUsagePortion
	rejectDurableConsumers
	rollbackOnlyOnAsyncException
	schedulePeriodForDestinationPurge
	schedulerSupport
	splitSystemUsageForProducersConsumers
	taskRunnerPriority

Element	Attribute
	timeBeforePurgeTempDestinations
	useAuthenticatedPrincipalFo rJMSXUserID
	useMirroredQueues
	useTempMirroredQueues
	useVirtualDestSubs
	useVirtualDestSubsOnCreation
	useVirtualTopics
cachedMessageGroupMapFactory	cacheSize
compositeQueue	concurrentSend
	copyMessage
	forwardOnly
	name
	sendWhenNotMatched
compositeTopic	concurrentSend
	copyMessage
	forwardOnly
	name
	sendWhenNotMatched
conditionalNetworkBridgeFilterFactory	rateDuration
	rateLimit

Element	Attribute
	`replayDelay`
	`replayWhenNoConsumers`
	`selectorAware` **Supported in** Apache ActiveMQ 5.16.x
`constantPendingMessageLimitStrategy`	`limit`
`discarding`	`deadLetterQueue`
	`enableAudit`
	`expiration`
	`maxAuditDepth`
	`maxProducersToAudit`
	`processExpired`
	`processNonPersistent`
`discardingDLQBrokerPlugin`	`dropAll`
	`dropOnly`
	`dropTemporaryQueues`
	`dropTemporaryTopics`
	`reportInterval`
`filteredDestination`	`queue`

Element	Attribute
	selector
	topic
fixedCountSubscriptionRecoveryPolicy	maximumSize
fixedSizedSubscriptionRecoveryPolicy	maximumSize
	useSharedBuffer
forcePersistencyModeBrokerPlugin	persistenceFlag
individualDeadLetterStrategy	destinationPerDurableSubscriber
	enableAudit
	expiration
	maxAuditDepth
	maxProducersToAudit
	processExpired
	processNonPersistent
	queuePrefix
	queueSuffix
	topicPrefix
	topicSuffix
	useQueueForQueueMessages
	useQueueForTopicMessages
messageGroupHashBucketFactory	bucketCount

Element	Attribute
	cacheSize
mirroredQueue	copyMessage
	postfix
	prefix
oldestMessageEvictionStrategy	evictExpiredMessagesHighWatermark
oldestMessageWithLowestPriorityEvictionStrategy	evictExpiredMessagesHighWatermark
policyEntry	advisoryForConsumed
	advisoryForDelivery
	advisoryForDiscardingMessages
	advisoryForFastProducers
	advisoryForSlowConsumers
	advisoryWhenFull
	allConsumersExclusiveByDefault
	alwaysRetroactive
	blockedProducerWarningInterval
	consumersBeforeDispatchStarts
	cursorMemoryHighWaterMark
	doOptimzeMessageStorage
	durableTopicPrefetch

Element	Attribute
	enableAudit
	expireMessagesPeriod
	gcInactiveDestinations
	gcWithNetworkConsumers
	inactiveTimeoutBeforeGC
	inactiveTimoutBeforeGC
	includeBodyForAdvisory
	lazyDispatch
	maxAuditDepth
	maxBrowsePageSize
	maxDestinations
	maxExpirePageSize
	maxPageSize
	maxProducersToAudit
	maxQueueAuditDepth
	memoryLimit
	messageGroupMapFactoryType
	minimumMessageSize
	optimizedDispatch
	optimizeMessageStoreInFligh tLimit

Element	Attribute
	persistJMSRedelivered
	prioritizedMessages
	producerFlowControl
	queue
	queueBrowserPrefetch
	queuePrefetch
	reduceMemoryFootprint
	sendAdvisoryIfNoConsumers
	sendFailIfNoSpace
	sendFailIfNoSpaceAfterTimeout **Supported in** Apache ActiveMQ 5.16.4 and above
	sendDuplicateFromStoreToDLQ
	storeUsageHighWaterMark
	strictOrderDispatch
	tempQueue
	tempTopic
	timeBeforeDispatchStarts
	topic
	topicPrefetch

Element	Attribute
	useCache
	useConsumerPriority
usePrefetchExtension	
prefetchRatePendingMessageLimitStrategy	multiplier
queryBasedSubscriptionRecoveryPolicy	query
queue	DLQ
	physicalName
redeliveryPlugin	fallbackToDeadLetter
	sendToDlqIfMaxRetriesExceeded
redeliveryPolicy	backOffMultiplier
	collisionAvoidancePercent
	initialRedeliveryDelay
	maximumRedeliveries
	maximumRedeliveryDelay
	preDispatchCheck
	queue
	redeliveryDelay
	tempQueue
	tempTopic

Element	Attribute
	topic
	useCollisionAvoidance
	useExponentialBackOff
sharedDeadLetterStrategy	enableAudit
	expiration
	maxAuditDepth
	maxProducersToAudit
	processExpired
	processNonPersistent
storeDurableSubscriberCursor	immediatePriorityDispatch
	useCache
tempDestinationAuthorizationEntry	admin
	queue
	read
	tempQueue
	tempTopic
	topic
	write
tempQueue	DLQ
	physicalName

Element	Attribute
tempTopic	DLQ
	physicalName
timedSubscriptionRecoveryPolicy	zeroExpirationOverride
timeStampingBrokerPlugin	recoverDuration
	futureOnly
	processNetworkMessages
	ttlCeiling
topic	DLQ
	physicalName
transportConnector	•
	name
	updateClusterClients
	rebalanceClusterClients
	updateClusterClientsOnRemove
uniquePropertyMessageEvictionStrategy	evictExpiredMessagesHighWatermark
	propertyName
virtualTopic	concurrentSend
	local
	dropOnResourceLimit

Element	Attribute
	name
	postfix
	prefix
	selectorAware
	setOriginalDestination
	transactedSend

Amazon MQ Parent Element Attributes

The following is a detailed explanation of parent element attributes. For more information, see XML Configuration in the Apache ActiveMQ documentation.

Topics

- broker

broker

`broker` is a parent collection element.

Attributes

networkConnectionStartAsync

To mitigate network latency and to allow other networks to start in a timely manner, use the `<networkConnectionStartAsync>` tag. The tag instructs the broker to use an executor to start network connections in parallel, asynchronous to a broker start.

Default: `false`

Example Configuration

```
<broker networkConnectorStartAsync="false"/>
```

Elements, Child Collection Elements, and Their Child Elements Permitted in Amazon MQ Configurations

The following is a detailed listing of the elements, child collection elements, and their child elements permitted in Amazon MQ configurations. For more information, see XML Configuration in the Apache ActiveMQ documentation.

Element	Child Collection Element	Child Element
authorizationMap	authorizationEntries	authorizationEntry
		tempDestinationAuthorizationEntry
	defaultEntry	authorizationEntry
		tempDestinationAuthorizationEntry
	tempDestinationAuthorizationEntry	tempDestinationAuthorizationEntry
authorizationPlugin	map	authorizationMap
broker	destinationInterceptors	mirroredQueue
		virtualDestinationInterceptor
	destinationPolicy	policyMap
	destinations	queue
		tempQueue
		tempTopic
		topic
	networkConnectors	networkConnector

Element	Child Collection Element	Child Element
	persistenceAdapter	kahaDB
	plugins	authorizationPlugin
		discardingDLQBrokerPlugin
		forcePersistencyModeBrokerPlugin
		redeliveryPlugin
		statisticsBrokerPlugin
		timeStampingBrokerPlugin
	systemUsage	systemUsage
	transportConnector	name
		updateClusterClients
		rebalanceClusterClients
		updateClusterClientsOnRemove
compositeQueue	forwardTo	queue
		tempQueue
		tempTopic
		topic
		filteredDestination

Element	Child Collection Element	Child Element
compositeTopic	forwardTo	queue
		tempQueue
		tempTopic
		topic
		filteredDestination
policyEntry	deadLetterStrategy	discarding
		individualDeadLetterStrategy
		sharedDeadLetterStrategy
	destination	queue
		tempQueue
		tempTopic
		topic
	dispatchPolicy	priorityDispatchPolicy
		priorityNetworkDispatchPolicy
		roundRobinDispatchPolicy
		simpleDispatchPolicy
		strictOrderDispatchPolicy

Element	Child Collection Element	Child Element
		clientIdFilterDispatchPolicy
	messageEvictionStrategy	oldestMessageEvictionStrategy
		oldestMessageWithLowestPriorityEvictionStrategy
		uniquePropertyMessageEvictionStrategy
	messageGroupMapFactory	cachedMessageGroupMapFactory
		messageGroupHashBucketFactory
		simpleMessageGroupMapFactory
	pendingDurableSubscriberPolicy	fileDurableSubscriberCursor
		storeDurableSubscriberCursor
		vmDurableCursor
	pendingMessageLimitStrategy	constantPendingMessageLimitStrategy
		prefetchRatePendingMessageLimitStrategy

Element	Child Collection Element	Child Element
	pendingQueuePolicy	fileQueueCursor
		storeCursor
		vmQueueCursor
	pendingSubscriberPolicy	fileCursor
		vmCursor
	slowConsumerStrategy	abortSlowAckConsumerStrategy
		abortSlowConsumerStrategy
	subscriptionRecoveryPolicy	fixedCountSubscriptionRecoveryPolicy
		fixedSizedSubscriptionRecoveryPolicy
		lastImageSubscriptionRecoveryPolicy
		noSubscriptionRecoveryPolicy
		queryBasedSubscriptionRecoveryPolicy
		retainedMessageSubscriptionRecoveryPolicy
timedSubscriptionRecoveryPolicy		

Element	Child Collection Element	Child Element
policyMap	defaultEntry	policyEntry
	policyEntries	policyEntry
redeliveryPlugin	redeliveryPolicyMap	redeliveryPolicyMap
redeliveryPolicyMap	defaultEntry	redeliveryPolicy
	redeliveryPolicyEntries	redeliveryPolicy
retainedMessageSubscriptionRecoveryPolicy	wrapped	fixedCountSubscriptionRecoveryPolicy
		fixedSizedSubscriptionRecoveryPolicy
		lastImageSubscriptionRecoveryPolicy
		noSubscriptionRecoveryPolicy
		queryBasedSubscriptionRecoveryPolicy
		retainedMessageSubscriptionRecoveryPolicy
		timedSubscriptionRecoveryPolicy
sharedDeadLetterStrategy	deadLetterQueue	queue
		tempQueue
		tempTopic

Element	Child Collection Element	Child Element
		`topic`
`virtualDestinationInterceptor`	`virtualDestinations`	`compositeQueue`
		`compositeTopic`
		`virtualTopic`

Amazon MQ Child Element Attributes

The following is a detailed explanation of child element attributes. For more information, see XML Configuration in the Apache ActiveMQ documentation.

Topics

- authorizationEntry
- networkConnector
- kahaDB
- systemUsage

authorizationEntry

`authorizationEntry` is a child of the `authorizationEntries` child collection element.

Attributes

admin|read|write

The permissions granted to a group of users. For more information, see Always configure an authorization map.

If you specify an authorization map which doesn't include the `activemq-webconsole` group, you can't use the ActiveMQ Web Console because the group isn't authorized to send messages to, or receive messages from, the Amazon MQ broker.

Default: `null`

Example Configuration

```
<authorizationPlugin>
   <map>
      <authorizationMap>
         <authorizationEntries>
            <authorizationEntry admin="admins,activemq-webconsole"
 read="admins,users,activemq-webconsole" write="admins,activemq-webconsole" queue=">"/>
            <authorizationEntry admin="admins,activemq-webconsole"
 read="admins,users,activemq-webconsole" write="admins,activemq-webconsole" topic=">"/>
         </authorizationEntries>
      </authorizationMap>
   </map>
</authorizationPlugin>
```

networkConnector

`networkConnector` is a child of the `networkConnectors` child collection element.

Topics

- Attributes
- Example Configurations

Attributes

conduitSubscriptions

Specifies whether a network connection in a network of brokers treats multiple consumers subscribed to the same destination as one consumer. For example, if `conduitSubscriptions` is set to `true` and two consumers connect to broker B and consume from a destination, broker B combines the subscriptions into a single logical subscription over the network connection to broker A, so that only a single copy of a message is forwarded from broker A to broker B.

> **Note**
>
> Setting `conduitSubscriptions` to `true` can reduce redundant network traffic. However, using this attribute can have implications for the load-balancing of messages across consumers and might cause incorrect behavior in certain scenarios (for example, with JMS message selectors or with durable topics).

Default: `true`

duplex

Specifies whether the connection in the network of brokers is used to produce *and* consume messages. For example, if broker A creates a connection to broker B in non-duplex mode, messages can be forwarded only from broker A to broker B. However, if broker A creates a duplex connection to broker B, then broker B can forward messages to broker A without having to configure a `<networkConnector>`.

Default: `false`

name

The name of the bridge in the network of brokers.

Default: `bridge`

uri

The wire-level protocol endpoint for one of two brokers (or for multiple brokers) in a network of brokers.

Default: `null`

username

The username common to the brokers in a network of brokers.

Default: `null`

Example Configurations

> **Note**
>
> When using a `networkConnector` to define a network of brokers, don't include the password for the user common to your brokers.

A Network of Brokers with Two Brokers

In this configuration, two brokers are connected in a network of brokers. The name of the network connector is `connector_1_to_2`, the username common to the brokers is `myCommonUser`, the

connection is `duplex`, and the OpenWire endpoint URI is prefixed by `static:`, indicating a one-to-one connection between the brokers.

```
<networkConnectors>
  <networkConnector name="connector_1_to_2" userName="myCommonUser" duplex="true"
    uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

For more information, see Configure Network Connectors for Your Broker.

A Network of Brokers with Multiple Brokers

In this configuration, multiple brokers are connected in a network of brokers. The name of the network connector is `connector_1_to_2`, the username common to the brokers is `myCommonUser`, the connection is `duplex`, and the comma-separated list of OpenWire endpoint URIs is prefixed by `masterslave:`, indicating a failover connection between the brokers. The failover from broker to broker isn't randomized and reconnection attempts continue indefinitely.

```
<networkConnectors>
  <networkConnector name="connector_1_to_2" userName="myCommonUser" duplex="true"
    uri="masterslave:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617,
    ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-east-2.amazonaws.com:61617)"/>
</networkConnectors>
```

> **Note**
>
> We recommend using the `masterslave:` prefix for networks of brokers. The prefix is identical to the more explicit `static:failover:()?randomize=false&maxReconnectAttempts=0` syntax.

> **Note**
>
> This XML configuration does not allow spaces.

kahaDB

`kahaDB` is a child of the `persistenceAdapter` child collection element.

Attributes

concurrentStoreAndDispatchQueues

Specifies whether to use concurrent store and dispatch for queues. For more information, see Disable Concurrent Store and Dispatch for Queues with Slow Consumers.

Default: `true`

cleanupOnStop

> **Supported in**
>
> Apache ActiveMQ 15.16.x and above

If deactivated, garbage collection and cleanup does not take place when the broker is stopped, which speeds up the shutdown process. The increased speed is useful in cases with large databases or scheduler databases.

Default: `true`

journalDiskSyncInterval

Interval (ms) for when to perform a disk sync if `journalDiskSyncStrategy=periodic`. For more information, see the Apache ActiveMQ kahaDB documentation.

Default: `1000`

journalDiskSyncStrategy

> **Supported in**
>
> Apache ActiveMQ 15.14.x and above

Configures the disk sync policy. For more information, see the Apache ActiveMQ kahaDB documentation.

Default: `always`

> **Note**
>
> The ActiveMQ documentation states that the data loss is limited to the duration of `journalDiskSyncInterval`, which has a default of 1s. The data loss can be longer than the interval, but it is difficult to be precise. Use caution.

preallocationStrategy

Configures how the broker will try to preallocate the journal files when a new journal file is needed. For more information, see the Apache ActiveMQ kahaDB documentation.

Default: `sparse_file`

Example Configuration

Example

```
<broker xmlns="http://activemq.apache.org/schema/core">
    <persistenceAdapter>
       <kahaDB preallocationStrategy="zeros" concurrentStoreAndDispatchQueues="false"
 journalDiskSyncInterval="10000" journalDiskSyncStrategy="periodic"/>
   </persistenceAdapter>
</broker>
```

systemUsage

`systemUsage` is a child of the `systemUsage` child collection element. It controls the maximum amount of space the broker will use before slowing down producers. For more information, see Producer Flow Control in the Apache ActiveMQ documentation.

Child Element

memoryUsage

`memoryUsage` is a child of the `systemUsage` child element. It manages memory usage. Use `memoryUsage` to keep track of how much of something is being used so that you can control working set usage productively. For more information, see the schema in the Apache ActiveMQ documentation.

Child Element

`memoryUsage` is a child of the `memoryUsage` child element.

Attribute

percentOfJvmHeap

Integer between 0 (inclusive) and 70 (inclusive).

Default: `70`

Attributes

sendFailIfNoSpace

Sets whether a `send()` method should fail if there is no space free. The default value is false, which blocks the `send()` method until space becomes available. For more information, see the schema in the Apache Active MQ documentation.

Default: `false`

sendFailIfNoSpaceAfterTimeout

Default: `null`

Example Configuration

Example

```
<broker xmlns="http://activemq.apache.org/schema/core">
    <systemUsage>
      <systemUsage sendFailIfNoSpace="true" sendFailIfNoSpaceAfterTimeout="2000">
          <memoryUsage>
              <memoryUsage  percentOfJvmHeap="60" />
          </memoryUsage>>
      </systemUsage>
    </systemUsage>
</broker>
</persistenceAdapter>
```

Managing Amazon MQ for ActiveMQ engine versions

Apache ActiveMQ organizes version numbers according to semantic versioning specification as `X.Y.Z`. In Amazon MQ for ActiveMQ implementations, `X.Y` denotes the major version, and `Z` represents the minor version number. Amazon MQ considers a version change to be major if the major version numbers change. For example, upgrading from version 5.**15** to 5.**16** is considered a *major version upgrade*. A version change is considered minor if only the minor version number changes. For example, upgrading from version 5.15.**14** to 5.15.**15** is considered a *minor version upgrade*.

Amazon MQ for ActiveMQ currently supports the following engine versions of Apache ActiveMQ.

Major versions	Minor versions
ActiveMQ 5.17	• 5.17.6 (recommended)
ActiveMQ 5.16	• 5.16.7
ActiveMQ 5.15	• 5.15.16

The following minor versions have been deprecated.

Deprecated minor versions

Major versions	Minor versions
ActiveMQ 5.17	• 5.17.3 • 5.17.2 • 5.17.1
ActiveMQ 5.16	• 5.16.5 • 5.16.4 • 5.16.3 • 5.16.2
ActiveMQ 5.15	• 5.15.15 • 5.15.14

Major versions	Minor versions
	• 5.15.13 • 5.15.12 • 5.15.10 • 5.15.9 • 5.15.8 • 5.15.6 • 5.15.0

When you create a new Amazon MQ for ActiveMQ broker, you can specify any supported ActiveMQ engine version. If you use the AWS Management Console to create a broker, Amazon MQ automatically defaults to the latest engine version number. If you use the AWS CLI or the Amazon MQ API to create a broker, the engine version number is required. If you don't provide a version number, the operation will result in an exception. To learn more, see `create-broker` in the *AWS CLI Command Reference* and `CreateBroker` in the *Amazon MQ REST API Reference*.

Topics

Major and minor version upgrades

With Amazon MQ, you control when to upgrade your brokers to new versions. When automatic minor version upgrade is activated, Amazon MQ will automatically upgrade your broker engine to new minor versions as they are released and supported by Amazon MQ.

To perform a major version upgrade, you must manually upgrade your broker's engine version number. Minor and major version upgrades occure at the same time as other broker patching operations, during your scheduled maintenance window. If you opt out of automatic minor version upgrades, you can manually upgrade your broker to a new supported minor version by following the same procedure as a major upgrade.

For more information about updating your broker preferences to activate or deactivate minor version upgrades, and manually upgrading your broker, see the section called “Upgrading the engine version”.

Listing supported engine versions

You can list all supported minor and major engine versions by using the `describe-broker-instance-options` AWS CLI command.

```
aws mq describe-broker-instance-options
```

To filter the results by engine and instance type use the `--engine-type` and `--host-instance-type` options as shown in the following.

```
aws mq describe-broker-instance-options --engine-type engine-type --host-instance-
type instance-type
```

For example, to filter the results for ActiveMQ, and `mq.m5.large` instance type, replace *`engine-type`* with `ACTIVEMQ` and *`instance-type`* with `mq.m5.large`.

Working examples of using Java Message Service (JMS) with ActiveMQ

The following examples show how you can work with ActiveMQ programmatically:

- The OpenWire example Java code connects to a broker, creates a queue, and sends and receives a message. For a detailed breakdown and explanation, see Connecting a Java application to your broker.
- The MQTT example Java code connects to a broker, creates a topic, and publishes and receives a message.
- The STOMP+WSS example Java code connects to a broker, creates a queue, and publishes and receives a message.

Prerequisites

Enable VPC Attributes

To ensure that your broker is accessible within your VPC, you must enable the `enableDnsHostnames` and `enableDnsSupport` VPC attributes. For more information, see DNS Support in your VPC in the *Amazon VPC User Guide*.

Enable inbound Connections

1. Sign in to the Amazon MQ console.

2. From the broker list, choose the name of your broker (for example, **MyBroker**).
3. On the ***MyBroker*** page, in the **Connections** section, note the addresses and ports of the broker's web console URL and wire-level protocols.
4. In the **Details** section, under **Security and network**, choose the name of your security group or .

   The **Security Groups** page of the EC2 Dashboard is displayed.
5. From the security group list, choose your security group.
6. At the bottom of the page, choose **Inbound**, and then choose **Edit**.
7. In the **Edit inbound rules** dialog box, add a rule for every URL or endpoint that you want to be publicly accessible (the following example shows how to do this for a broker web console).

   a. Choose **Add Rule**.

   b. For **Type**, select **Custom TCP**.

   c. For **Port Range**, type the web console port (8162).

   d. For **Source**, leave **Custom** selected and then type the IP address of the system that you want to be able to access the web console (for example, `192.0.2.1`).

   e. Choose **Save**.

   Your broker can now accept inbound connections.

Add Java dependencies

OpenWire

Add the `activemq-client.jar` and `activemq-pool.jar` packages to your Java class path. The following example shows these dependencies in a Maven project `pom.xml` file.

```
<dependencies>
    <dependency>
        <groupId>org.apache.activemq</groupId>
        <artifactId>activemq-client</artifactId>
        <version>5.15.16</version>
    </dependency>
    <dependency>
        <groupId>org.apache.activemq</groupId>
        <artifactId>activemq-pool</artifactId>
```

```
        <version>5.15.16</version>
    </dependency>
</dependencies>
```

For more information about `activemq-client.jar`, see Initial Configuration in the Apache ActiveMQ documentation.

MQTT

Add the `org.eclipse.paho.client.mqttv3.jar` package to your Java class path. The following example shows this dependency in a Maven project `pom.xml` file.

```
<dependencies>
    <dependency>
        <groupId>org.eclipse.paho</groupId>
        <artifactId>org.eclipse.paho.client.mqttv3</artifactId>
        <version>1.2.0</version>
    </dependency>
    </dependencies>
```

For more information about `org.eclipse.paho.client.mqttv3.jar`, see Eclipse Paho Java Client.

STOMP+WSS

Add the following packages to your Java class path:

- `spring-messaging.jar`
- `spring-websocket.jar`
- `javax.websocket-api.jar`
- `jetty-all.jar`
- `slf4j-simple.jar`
- `jackson-databind.jar`

The following example shows these dependencies in a Maven project `pom.xml` file.

```
<dependencies>
    <dependency>
        <groupId>org.springframework</groupId>
```

```
        <artifactId>spring-messaging</artifactId>
        <version>5.0.5.RELEASE</version>
    </dependency>
    <dependency>
        <groupId>org.springframework</groupId>
        <artifactId>spring-websocket</artifactId>
        <version>5.0.5.RELEASE</version>
    </dependency>
    <dependency>
        <groupId>javax.websocket</groupId>
        <artifactId>javax.websocket-api</artifactId>
        <version>1.1</version>
    </dependency>
    <dependency>
        <groupId>org.eclipse.jetty.aggregate</groupId>
        <artifactId>jetty-all</artifactId>
        <type>pom</type>
        <version>9.3.3.v20150827</version>
    </dependency>
    <dependency>
        <groupId>org.slf4j</groupId>
        <artifactId>slf4j-simple</artifactId>
        <version>1.6.6</version>
    </dependency>
    <dependency>
        <groupId>com.fasterxml.jackson.core</groupId>
        <artifactId>jackson-databind</artifactId>
        <version>2.5.0</version>
    </dependency>
    </dependencies>
```

For more information, see STOMP Support in the Spring Framework documentation.

AmazonMQExample.java

Important

In the following example code, producers and consumers run in a single thread. For production systems (or to test broker instance failover), make sure that your producers and consumers run on separate hosts or threads.

OpenWire

```
/*
 * Copyright 2010-2019 Amazon.com, Inc. or its affiliates. All Rights Reserved.
 *
 * Licensed under the Apache License, Version 2.0 (the "License").
 * You may not use this file except in compliance with the License.
 * A copy of the License is located at
 *
 *  https://aws.amazon.com/apache2.0
 *
 * or in the "license" file accompanying this file. This file is distributed
 * on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either
 * express or implied. See the License for the specific language governing
 * permissions and limitations under the License.
 *
 */

    import org.apache.activemq.ActiveMQConnectionFactory;
    import org.apache.activemq.jms.pool.PooledConnectionFactory;

    import javax.jms.*;

    public class AmazonMQExample {

    // Specify the connection parameters.
    private final static String WIRE_LEVEL_ENDPOINT
            = "ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617";
    private final static String ACTIVE_MQ_USERNAME = "MyUsername123";
    private final static String ACTIVE_MQ_PASSWORD = "MyPassword456";

    public static void main(String[] args) throws JMSException {
        final ActiveMQConnectionFactory connectionFactory =
                createActiveMQConnectionFactory();
        final PooledConnectionFactory pooledConnectionFactory =
                createPooledConnectionFactory(connectionFactory);

        sendMessage(pooledConnectionFactory);
        receiveMessage(connectionFactory);

        pooledConnectionFactory.stop();
    }
```

```
    private static void
    sendMessage(PooledConnectionFactory pooledConnectionFactory) throws JMSException
{
        // Establish a connection for the producer.
        final Connection producerConnection = pooledConnectionFactory
                .createConnection();
        producerConnection.start();

        // Create a session.
        final Session producerSession = producerConnection
                .createSession(false, Session.AUTO_ACKNOWLEDGE);

        // Create a queue named "MyQueue".
        final Destination producerDestination = producerSession
                .createQueue("MyQueue");

        // Create a producer from the session to the queue.
        final MessageProducer producer = producerSession
                .createProducer(producerDestination);
        producer.setDeliveryMode(DeliveryMode.NON_PERSISTENT);

        // Create a message.
        final String text = "Hello from Amazon MQ!";
        final TextMessage producerMessage = producerSession
                .createTextMessage(text);

        // Send the message.
        producer.send(producerMessage);
        System.out.println("Message sent.");

        // Clean up the producer.
        producer.close();
        producerSession.close();
        producerConnection.close();
    }

    private static void
    receiveMessage(ActiveMQConnectionFactory connectionFactory) throws JMSException
{
        // Establish a connection for the consumer.
        // Note: Consumers should not use PooledConnectionFactory.
        final Connection consumerConnection = connectionFactory.createConnection();
        consumerConnection.start();
```

```
        // Create a session.
        final Session consumerSession = consumerConnection
                .createSession(false, Session.AUTO_ACKNOWLEDGE);

        // Create a queue named "MyQueue".
        final Destination consumerDestination = consumerSession
                .createQueue("MyQueue");

        // Create a message consumer from the session to the queue.
        final MessageConsumer consumer = consumerSession
                .createConsumer(consumerDestination);

        // Begin to wait for messages.
        final Message consumerMessage = consumer.receive(1000);

        // Receive the message when it arrives.
        final TextMessage consumerTextMessage = (TextMessage) consumerMessage;
        System.out.println("Message received: " + consumerTextMessage.getText());

        // Clean up the consumer.
        consumer.close();
        consumerSession.close();
        consumerConnection.close();
    }

    private static PooledConnectionFactory
    createPooledConnectionFactory(ActiveMQConnectionFactory connectionFactory) {
        // Create a pooled connection factory.
        final PooledConnectionFactory pooledConnectionFactory =
                new PooledConnectionFactory();
        pooledConnectionFactory.setConnectionFactory(connectionFactory);
        pooledConnectionFactory.setMaxConnections(10);
        return pooledConnectionFactory;
    }

    private static ActiveMQConnectionFactory createActiveMQConnectionFactory() {
        // Create a connection factory.
        final ActiveMQConnectionFactory connectionFactory =
                new ActiveMQConnectionFactory(WIRE_LEVEL_ENDPOINT);

        // Pass the sign-in credentials.
        connectionFactory.setUserName(ACTIVE_MQ_USERNAME);
        connectionFactory.setPassword(ACTIVE_MQ_PASSWORD);
        return connectionFactory;
```

```
    }
    }
```

MQTT

```
/*
 * Copyright 2010-2019 Amazon.com, Inc. or its affiliates. All Rights Reserved.
 *
 * Licensed under the Apache License, Version 2.0 (the "License").
 * You may not use this file except in compliance with the License.
 * A copy of the License is located at
 *
 *  https://aws.amazon.com/apache2.0
 *
 * or in the "license" file accompanying this file. This file is distributed
 * on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either
 * express or implied. See the License for the specific language governing
 * permissions and limitations under the License.
 *
 */

    import org.eclipse.paho.client.mqttv3.*;

    public class AmazonMQExampleMqtt implements MqttCallback {

    // Specify the connection parameters.
    private final static String WIRE_LEVEL_ENDPOINT =
            "ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:8883";
    private final static String ACTIVE_MQ_USERNAME = "MyUsername123";
    private final static String ACTIVE_MQ_PASSWORD = "MyPassword456";

    public static void main(String[] args) throws Exception {
        new AmazonMQExampleMqtt().run();
    }

    private void run() throws MqttException, InterruptedException {

        // Specify the topic name and the message text.
        final String topic = "myTopic";
        final String text = "Hello from Amazon MQ!";

        // Create the MQTT client and specify the connection options.
```

```
        final String clientId = "abc123";
        final MqttClient client = new MqttClient(WIRE_LEVEL_ENDPOINT, clientId);
        final MqttConnectOptions connOpts = new MqttConnectOptions();

        // Pass the sign-in credentials.
        connOpts.setUserName(ACTIVE_MQ_USERNAME);
        connOpts.setPassword(ACTIVE_MQ_PASSWORD.toCharArray());

        // Create a session and subscribe to a topic filter.
        client.connect(connOpts);
        client.setCallback(this);
        client.subscribe("+");

        // Create a message.
        final MqttMessage message = new MqttMessage(text.getBytes());

        // Publish the message to a topic.
        client.publish(topic, message);
        System.out.println("Published message.");

        // Wait for the message to be received.
        Thread.sleep(3000L);

        // Clean up the connection.
        client.disconnect();
    }

    @Override
    public void connectionLost(Throwable cause) {
        System.out.println("Lost connection.");
    }

    @Override
    public void messageArrived(String topic, MqttMessage message) throws
MqttException {
        System.out.println("Received message from topic " + topic + ": " + message);
    }

    @Override
    public void deliveryComplete(IMqttDeliveryToken token) {
        System.out.println("Delivered message.");
    }
    }
```

STOMP+WSS

```
/*
 * Copyright 2010-2019 Amazon.com, Inc. or its affiliates. All Rights Reserved.
 *
 * Licensed under the Apache License, Version 2.0 (the "License").
 * You may not use this file except in compliance with the License.
 * A copy of the License is located at
 *
 *  https://aws.amazon.com/apache2.0
 *
 * or in the "license" file accompanying this file. This file is distributed
 * on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either
 * express or implied. See the License for the specific language governing
 * permissions and limitations under the License.
 *
 */

    import org.springframework.messaging.converter.StringMessageConverter;
    import org.springframework.messaging.simp.stomp.*;
    import org.springframework.web.socket.WebSocketHttpHeaders;
    import org.springframework.web.socket.client.WebSocketClient;
    import org.springframework.web.socket.client.standard.StandardWebSocketClient;
    import org.springframework.web.socket.messaging.WebSocketStompClient;

    import java.lang.reflect.Type;

    public class AmazonMQExampleStompWss {

    // Specify the connection parameters.
    private final static String DESTINATION = "/queue";
    private final static String WIRE_LEVEL_ENDPOINT =
            "wss://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61619";
    private final static String ACTIVE_MQ_USERNAME = "MyUsername123";
    private final static String ACTIVE_MQ_PASSWORD = "MyPassword456";

    public static void main(String[] args) throws Exception {
        final AmazonMQExampleStompWss example = new AmazonMQExampleStompWss();

        final StompSession stompSession = example.connect();
        System.out.println("Subscribed to a destination using session.");
        example.subscribeToDestination(stompSession);
```

```
        System.out.println("Sent message to session.");
        example.sendMessage(stompSession);
        Thread.sleep(60000);
    }

    private StompSession connect() throws Exception {
        // Create a client.
        final WebSocketClient client = new StandardWebSocketClient();
        final WebSocketStompClient stompClient = new WebSocketStompClient(client);
        stompClient.setMessageConverter(new StringMessageConverter());

        final WebSocketHttpHeaders headers = new WebSocketHttpHeaders();

        // Create headers with authentication parameters.
        final StompHeaders head = new StompHeaders();
        head.add(StompHeaders.LOGIN, ACTIVE_MQ_USERNAME);
        head.add(StompHeaders.PASSCODE, ACTIVE_MQ_PASSWORD);

        final StompSessionHandler sessionHandler = new MySessionHandler();

        // Create a connection.
        return stompClient.connect(WIRE_LEVEL_ENDPOINT, headers, head,
                sessionHandler).get();
    }

    private void subscribeToDestination(final StompSession stompSession) {
        stompSession.subscribe(DESTINATION, new MyFrameHandler());
    }

    private void sendMessage(final StompSession stompSession) {
        stompSession.send(DESTINATION, "Hello from Amazon MQ!".getBytes());
    }

    private static class MySessionHandler extends StompSessionHandlerAdapter {
        public void afterConnected(final StompSession stompSession,
                                   final StompHeaders stompHeaders) {
            System.out.println("Connected to broker.");
        }
    }

    private static class MyFrameHandler implements StompFrameHandler {
        public Type getPayloadType(final StompHeaders headers) {
            return String.class;
        }
```

```
        public void handleFrame(final StompHeaders stompHeaders,
                                final Object message) {
            System.out.print("Received message from topic: " + message);
        }
    }
    }
```

ActiveMQ tutorials

The following tutorials show how you can create and connect to your ActiveMQ brokers. To use the ActiveMQ Java example code, you must install the Java Standard Edition Development Kit and make some changes to the code

Topics

- Creating and configuring an ActiveMQ broker
- Creating and configuring an Amazon MQ network of brokers
- Connecting a Java application to your Amazon MQ broker
- Integrating ActiveMQ brokers with LDAP
- Creating and managing ActiveMQ broker users

Creating and configuring an ActiveMQ broker

A *broker* is a message broker environment running on Amazon MQ. It is the basic building block of Amazon MQ. The combined description of the broker instance *class* (`m5`, `t3`) and *size* (`large`, `micro`) is a *broker instance type* (for example, `mq.m5.large`). For more information, see Broker.

The first and most common Amazon MQ task is creating a broker. The following example shows how you can create and configure a broker using the AWS Management Console.

Topics

- Step 1: Configure Basic Broker Settings
- Step 2: (Optional) Configure Additional Broker Settings
- Step 3: Finish Creating the Broker
- Editing broker engine version, instance type, CloudWatch logs, and maintenance preferences

Step 1: Configure Basic Broker Settings

1. Sign in to the Amazon MQ console.
2. On the **Select broker engine** page, choose **Apache ActiveMQ**.
3. On the **Select deployment and storage** page, in the **Deployment mode and storage type** section, do the following:

   a. Choose the **Deployment mode** (for example, **Active/standby broker**). For more information, see Broker Architecture.

      - A **Single-instance broker** is comprised of one broker in one Availability Zone. The broker communicates with your application and with an Amazon EBS or Amazon EFS storage volume. For more information, see Amazon MQ single-instance broker.
      - An **Active/standby broker for high availability** is comprised of two brokers in two different Availability Zones, configured in a *redundant pair*. These brokers communicate synchronously with your application, and with Amazon EFS. For more information, see Amazon MQ active/standby broker for high availability.
      - For more information on the sample blueprints for a network of brokers, see Sample blueprints.

   b. Choose the **Storage type** (for example, **EBS**). For more information, see Storage.

      > **Note**
      >
      > Amazon EBS replicates data within a single Availability Zone and doesn't support the ActiveMQ active/standby deployment mode.

   c. Choose **Next**.
4. On the **Configure settings** page, in the **Details** section, do the following:

   a. Enter the **Broker name**.

      > **Important**
      >
      > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker names. Broker names are accessible to other AWS services, including CloudWatch Logs. Broker names are not intended to be used for private or sensitive data.

b. Choose the **Broker instance type** (for example, **mq.m5.large**). For more information, see Broker instance types.

5. In the **ActiveMQ Web Console access** section, provide a **Username** and **Password**. The following restrictions apply to broker usernames and passwords:

- Your username can contain only alphanumeric characters, dashes, periods, underscores, and tildas (- . _ ~).
- Your password must be at least 12 characters long, contain at least 4 unique characters and must not contain commas, colons, or equal signs (,:=).

> **Important**
>
> Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

Step 2: (Optional) Configure Additional Broker Settings

> **Important**
>
> - **Subnet(s)** – A single-instance broker requires one subnet (for example, the default subnet). An active/standby broker requires two subnets.
> - **Security group(s)** – Both single-instance brokers and active/standby brokers require at least one security group (for example, the default security group).
> - **VPC** – A broker's subnet(s) and security group(s) must be in the same VPC. EC2-Classic resources aren't supported. Amazon MQ only supports default VPC tenancy, and does not support dedicated VPC tenancy.
> - **Encryption** – Choose the customer master key to encrypt your data. See Encryption at rest.
> - **Public accessibility** – Disabling public accessibility makes the broker accessible only within your VPC. For more information, see Prefer brokers without public accessibility and Accessing the broker web console without public accessibility.

1. Expand the **Additional settings** section.
2. In the **Configuration** section, choose **Create a new configuration with default values** or **Select an existing configuration**. For more information, see Configuration and Amazon MQ Broker Configuration Parameters.
3. In the **Logs** section, choose whether to publish **General** logs and **Audit** logs to Amazon CloudWatch Logs. For more information, see Configuring Amazon MQ to publish logs to Amazon CloudWatch Logs.

   > **Important**
   >
   > If you don't add the `CreateLogGroup` permission to your Amazon MQ user before the user creates or reboots the broker, Amazon MQ doesn't create the log group.
   > If you don't configure a resource-based policy for Amazon MQ, the broker can't publish the logs to CloudWatch Logs.

4. In the **Network and security section**, configure your broker's connectivity:
   a. Do one of the following:
      - Choose **Use the default VPC, subnet(s), and security group(s).**
      - Choose **Select existing VPC, subnet(s), and security group(s).**
        1. If you choose this option, you can create a new **Virtual Private Cloud (VPC)** on the Amazon VPC console, select an existing VPC, or select the default VPC. For more information, see What is Amazon VPC? in the *Amazon VPC User Guide*.
        2. After you create or select a VPC, you can create new **Subnet(s)** on the Amazon VPC console or select existing ones. For more information, see VPCs and Subnets in the *Amazon VPC User Guide*.
        3. After you create or select subnets, you can select the **Security group(s)**.
   b. Choose the customer master key (CMK) that will be used to encrypt your data. See Encryption at rest.
   c. Choose the **Public accessibility** of your broker.
5. In the **Maintenance section**, configure your broker's maintenance schedule:
   a. To upgrade the broker to new versions as Apache releases them, choose **Enable automatic minor version upgrades**. Automatic upgrades occur during the *maintenance window*

defined by the day of the week, the time of day (in 24-hour format), and the time zone (UTC by default).

> **Note**
>
> For an active/standby broker, if one of the broker instances undergoes maintenance, it takes Amazon MQ a short while to take the inactive instance out of service. This allows the healthy standby instance to become active and to begin accepting incoming communications.

b. Do one of the following:

- To allow Amazon MQ to select the maintenance window automatically, choose **No preference**.
- To set a custom maintenance window, choose **Select maintenance window** and then specify the **Start day** and **Start time** of the upgrades.

Step 3: Finish Creating the Broker

1. Choose **Deploy**.

   While Amazon MQ creates your broker, it displays the **Creation in progress** status.

   Creating the broker takes about 15 minutes.

   When your broker is created successfully, Amazon MQ displays the **Running** status.

	Name	Status	Deployment mode	Instance type
○	MyBroker	Running	Single-instance broker	mq.m5.large

2. Choose ***MyBroker***.

   On the ***MyBroker*** page, in the **Connect** section, note your broker's **ActiveMQ web console** URL, for example:

```
https://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:8162
```

Also, note your broker's wire-level protocol **Endpoints**. The following is an example of an OpenWire endpoint:

```
ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61617
```

> **Note**
>
> For an active/standby broker, Amazon MQ provides two ActiveMQ Web Console URLs, but only one URL is active at a time. Likewise, Amazon MQ provides two endpoints for each wire-level protocol, but only one endpoint is active in each pair at a time. The `-1` and `-2` suffixes denote a redundant pair. For more information, see Broker Architecture).
> For wire-level protocol endpoints, you can allow your application to connect to either endpoint by using the Failover Transport.

Editing broker engine version, instance type, CloudWatch logs, and maintenance preferences

In addition to editing broker configurations and managing configuration revisions, you can configure preferences specific to the broker.

> **Note**
>
> All preferences except for those for automatic minor version upgrades require you to schedule modifications. For more information, see Amazon MQ broker configuration lifecycle.

The following example shows how you can edit Amazon MQ ActiveMQ broker preferences using the AWS Management Console.

Edit ActiveMQ broker options

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Edit**.

3. On the **Edit *MyBroker*** page, in the **Specifications** section, select a **Broker engine version** or a **Broker Instance type**.

4. In the **Configuration** section, select the configuration and revision for your broker. For more information, see Creating and applying broker configurations.
5. In the **Security and network** section, select a group from the **Security group(s)** drop-down, or choose **Create a new security group** to open the Amazon VPC console.
6. In the **CloudWatch Logs** section, choose whether to publish **General** logs and **Audit** logs to Amazon CloudWatch Logs.

   For more information about configuring CloudWatch logs for ActiveMQ brokers, see Configuring Amazon MQ to publish logs to Amazon CloudWatch Logs.

   > **Important**
   >
   > If you don't add the `CreateLogGroup` permission to your Amazon MQ user before the user creates or reboots the broker, Amazon MQ doesn't create the log group.
   > If you don't configure a resource-based policy for Amazon MQ, the broker can't publish the logs to CloudWatch Logs.

7. In the **Maintenance** section, configure your broker's maintenance schedule:

   To upgrade the broker to new versions as AWS releases them, choose **Enable automatic minor version upgrades**. Automatic upgrades occur during the *maintenance window* defined by the day of the week, the time of day (in 24-hour format), and the time zone (UTC by default).

   > **Note**
   >
   > For an active/standby broker, if one of the broker instances undergoes maintenance, it takes Amazon MQ a short while to take the inactive instance out of service. This allows the healthy standby instance to become active and to begin accepting incoming communications.

8. Choose **Schedule modifications**.

> **Note**
>
> If you choose only **Enable automatic minor version upgrades**, the button changes to **Save** because no broker reboot is necessary.

Your preferences are applied to your broker at the specified time.

Creating and configuring an Amazon MQ network of brokers

A *network of brokers* is comprised of multiple simultaneously active single-instance brokers or active/standby brokers. You can configure networks of brokers in a variety of topologies (for example, *concentrator*, *hub-and-spokes*, *tree*, or *mesh*), depending on your application's needs, such as high availability and scalability. For instance, a hub and spoke network of brokers can increase resiliency, preserving messages if one broker is not reachable. A network of brokers with a concentrator topology can collect messages from a larger number of brokers accepting incoming messages, and concentrate them to more central brokers, to better handle the load of many incoming messages. In this tutorial, you learn how to create a two-broker network of brokers with a *source and sink* topology.

For a conceptual overview and detailed configuration information, see the following:

- Amazon MQ Network of brokers
- Configure Your Network of Brokers Correctly
- `networkConnector`
- `networkConnectionStartAsync`
- Networks of Brokers in the ActiveMQ documentation

You can use the Amazon MQ console to create an Amazon MQ network of brokers. Because you can start the creation of the two brokers in parallel, this process takes approximately 15 minutes.

Topics

- Next Steps

Prerequisites

To create a network of brokers, you must have the following:

- Two or more simultaneously active brokers (named `MyBroker1` and `MyBroker2` in this tutorial). For more information about creating brokers, see Creating and configuring a broker.
- The two brokers must be in the same VPC or in peered VPCs. For more information about VPCs, see What is Amazon VPC? in the *Amazon VPC User Guide* and What is VPC Peering? in the *Amazon VPC Peering Guide*.

> **Important**
>
> If you don't have a default VPC, subnet(s), or security group, you must create them first. For more information, see the following in the *Amazon VPC User Guide*:
>
> - Creating a Default VPC
> - Creating a Default Subnet
> - Creating a Security Group

- Two users with identical sign-in credentials for both brokers. For more information about creating users, see Creating and managing ActiveMQ broker users.

> **Note**
>
> When integrating LDAP authentication with a network of brokers, make sure that the user exists both as an ActiveMQ brokers, as well as an LDAP user.

The following example uses two single-instance brokers. However, you can create networks of brokers using active/standby brokers or a combination of broker deployment modes.

Step 1: Allow Traffic between Brokers

After you create your brokers, you must allow traffic between them.

1. On the Amazon MQ console, on the **MyBroker2** page, in the **Details** section, under **Security and network**, choose the name of your security group or .

   The **Security Groups** page of the EC2 Dashboard is displayed.
2. From the security group list, choose your security group.
3. At the bottom of the page, choose **Inbound**, and then choose **Edit**.
4. In the **Edit inbound rules** dialog box, add a rule for the OpenWire endpoint.

   a. Choose **Add Rule**.

   b. For **Type**, select **Custom TCP**.

   c. For **Port Range**, type the OpenWire port (61617).

   d. Do one of the following:

      - If you want to restrict access to a particular IP address, for **Source**, leave **Custom** selected, and then enter the IP address of `MyBroker1`, followed by `/32`. (This converts the IP address to a valid CIDR record). For more information see Elastic Network Interfaces.

        > **Tip**
        >
        > To retrieve the IP address of `MyBroker1`, on the Amazon MQ console, choose the name of the broker and navigate to the **Details** section.

      - If all the brokers are private and belong to the same VPC, for **Source**, leave **Custom** selected and then type the ID of the security group you are editing.

        > **Note**
        >
        > For public brokers, you must restrict access using IP addresses.

   e. Choose **Save**.

   Your broker can now accept inbound connections.

Step 2: Configure Network Connectors for Your Broker

After you allow traffic between your brokers, you must configure network connectors for one of them.

1. Edit the configuration revision for broker `MyBroker1`.

   a. On the **MyBroker1** page, choose **Edit**.

   b. On the **Edit MyBroker1** page, in the **Configuration** section, choose **View**.

      The broker engine type and version that the configuration uses (for example, **Apache ActiveMQ 5.15.0**) are displayed.

   c. On the **Configuration details** tab, the configuration revision number, description, and broker configuration in XML format are displayed.

   d. Choose **Edit configuration**.

   e. At the bottom of the configuration file, uncomment the `<networkConnectors>` section and include the following information:

      - The `name` for the network connector.
      - The ActiveMQ Web Console `username` that is common to both brokers.
      - Enable `duplex` connections.
      - Do one of the following:
        - If you are connecting the broker to a single-instance broker, use the `static:` prefix and the OpenWire endpoint `uri` for `MyBroker2`. For example:

          ```
          <networkConnectors>
            <networkConnector name="connector_1_to_2" userName="myCommonUser"
           duplex="true"
              uri="static:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
          east-2.amazonaws.com:61617)"/>
          </networkConnectors>
          ```

        - If you are connecting the broker to an active/standby broker, use the `static+failover` transport and the OpenWire endpoint `uri` for both brokers with the following query parameters `?randomize=false&maxReconnectAttempts=0`. For example:

          ```
          <networkConnectors>
          ```

```
  <networkConnector name="connector_1_to_2" userName="myCommonUser"
 duplex="true"
    uri="static:(failover:(ssl://
b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61617,
    ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)?randomize=false&amp;maxReconnectAttempts=0)"/>
</networkConnectors>
```

> **Note**
>
> Don't include the sign-in credentials for the ActiveMQ user.

f. Choose **Save**.

g. In the **Save revision** dialog box, type `Add network of brokers connector for MyBroker2`.

h. Choose **Save** to save the new revision of the configuration.

2. Edit `MyBroker1` to set the latest configuration revision to apply immediately.

   a. On the **MyBroker1** page, choose **Edit**.

   b. On the **Edit MyBroker1** page, in the **Configuration** section, choose **Schedule Modifications**.

   c. In the **Schedule broker modifications** section, choose to apply modifications **Immediately**.

   d. Choose **Apply**.

      `MyBroker1` is rebooted and your configuration revision is applied.

   The network of brokers is created.

Next Steps

After you configure your network of brokers, you can test it by producing and consuming messages.

> **Important**
>
> Make sure that you enable inbound connections *from your local machine* for broker `MyBroker1` on port 8162 (for the ActiveMQ Web Console) and port 61617 (for the OpenWire endpoint).
> You might also need to adjust your security group(s) settings to allow the producer and consumer to connect to the network of brokers.

1. On the Amazon MQ console, navigate to the **Connections** section and note the ActiveMQ Web Console endpoint for broker `MyBroker1`.
2. Navigate to the ActiveMQ Web Console for broker `MyBroker1`.
3. To verify that the network bridge is connected, choose **Network**.

   In the **Network Bridges** section, the name and the address of `MyBroker2` are listed in the **Remote Broker** and **Remote Address** columns.
4. From any machine that has access to broker `MyBroker2`, create a consumer. For example:

   ```
   activemq consumer --brokerUrl "ssl://
   b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com:61617" \
    --user commonUser \
    --password myPassword456 \
    --destination queue://MyQueue
   ```

   The consumer connects to the OpenWire endpoint of `MyBroker2` and begins to consume messages from queue `MyQueue`.
5. From any machine that has access to broker `MyBroker1`, create a producer and send some messages. For example:

   ```
   activemq producer --brokerUrl "ssl://
   b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-1.mq.us-east-2.amazonaws.com:61617" \
    --user commonUser \
    --password myPassword456 \
    --destination queue://MyQueue \
    --persistent true \
    --messageSize 1000 \
    --messageCount 10000
   ```

The producer connects to the OpenWire endpoint of `MyBroker1` and begins to produce persistent messages to queue `MyQueue`.

Connecting a Java application to your Amazon MQ broker

After you create an Amazon MQ ActiveMQ broker, you can connect your application to it. The following examples show how you can use the Java Message Service (JMS) to create a connection to the broker, create a queue, and send a message. For a complete, working Java example, see Working Java Example.

You can connect to ActiveMQ brokers using various ActiveMQ clients. We recommend using the ActiveMQ Client.

Topics

- Prerequisites
- To Create a Message Producer and Send a Message
- To Create a Message Consumer and Receive the Message

Prerequisites

Enable VPC Attributes

To ensure that your broker is accessible within your VPC, you must enable the `enableDnsHostnames` and `enableDnsSupport` VPC attributes. For more information, see DNS Support in your VPC in the *Amazon VPC User Guide*.

Enable Inbound Connections

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**).
3. On the ***MyBroker*** page, in the **Connections** section, note the addresses and ports of the broker's web console URL and wire-level protocols.
4. In the **Details** section, under **Security and network**, choose the name of your security group or .

   The **Security Groups** page of the EC2 Dashboard is displayed.

5. From the security group list, choose your security group.
6. At the bottom of the page, choose **Inbound**, and then choose **Edit**.
7. In the **Edit inbound rules** dialog box, add a rule for every URL or endpoint that you want to be publicly accessible (the following example shows how to do this for a broker web console).

   a. Choose **Add Rule**.

   b. For **Type**, select **Custom TCP**.

   c. For **Port Range**, type the web console port (`8162`).

   d. For **Source**, leave **Custom** selected and then type the IP address of the system that you want to be able to access the web console (for example, `192.0.2.1`).

   e. Choose **Save**.

   Your broker can now accept inbound connections.

Add Java Dependencies

Add the `activemq-client.jar` and `activemq-pool.jar` packages to your Java class path. The following example shows these dependencies in a Maven project `pom.xml` file.

```
<dependencies>
    <dependency>
        <groupId>org.apache.activemq</groupId>
        <artifactId>activemq-client</artifactId>
        <version>5.15.16</version>
    </dependency>
    <dependency>
        <groupId>org.apache.activemq</groupId>
        <artifactId>activemq-pool</artifactId>
        <version>5.15.16</version>
    </dependency>
</dependencies>
```

For more information about `activemq-client.jar`, see Initial Configuration in the Apache ActiveMQ documentation.

> **Important**
>
> In the following example code, producers and consumers run in a single thread. For production systems (or to test broker instance failover), make sure that your producers and consumers run on separate hosts or threads.

To Create a Message Producer and Send a Message

1. Create a JMS pooled connection factory for the message producer using your broker's endpoint and then call the `createConnection` method against the factory.

> **Note**
>
> For an active/standby broker, Amazon MQ provides two ActiveMQ Web Console URLs, but only one URL is active at a time. Likewise, Amazon MQ provides two endpoints for each wire-level protocol, but only one endpoint is active in each pair at a time. The `-1` and `-2` suffixes denote a redundant pair. For more information, see Broker Architecture).
> For wire-level protocol endpoints, you can allow your application to connect to either endpoint by using the Failover Transport.

```
// Create a connection factory.
final ActiveMQConnectionFactory connectionFactory = new
 ActiveMQConnectionFactory(wireLevelEndpoint);

// Pass the sign-in credentials.
connectionFactory.setUserName(activeMqUsername);
connectionFactory.setPassword(activeMqPassword);

// Create a pooled connection factory.
final PooledConnectionFactory pooledConnectionFactory = new
 PooledConnectionFactory();
pooledConnectionFactory.setConnectionFactory(connectionFactory);
pooledConnectionFactory.setMaxConnections(10);

// Establish a connection for the producer.
final Connection producerConnection = pooledConnectionFactory.createConnection();
```

```
producerConnection.start();

// Close all connections in the pool.
pooledConnectionFactory.clear();
```

> **Note**
>
> Message producers should always use the `PooledConnectionFactory` class. For more information, see Always Use Connection Pooling.

2. Create a session, a queue named `MyQueue`, and a message producer.

```
// Create a session.
final Session producerSession = producerConnection.createSession(false,
 Session.AUTO_ACKNOWLEDGE);

// Create a queue named "MyQueue".
final Destination producerDestination = producerSession.createQueue("MyQueue");

// Create a producer from the session to the queue.
final MessageProducer producer =
 producerSession.createProducer(producerDestination);
producer.setDeliveryMode(DeliveryMode.NON_PERSISTENT);
```

3. Create the message string `"Hello from Amazon MQ!"` and then send the message.

```
// Create a message.
final String text = "Hello from Amazon MQ!";
TextMessage producerMessage = producerSession.createTextMessage(text);

// Send the message.
producer.send(producerMessage);
System.out.println("Message sent.");
```

4. Clean up the producer.

```
producer.close();
producerSession.close();
producerConnection.close();
```

To Create a Message Consumer and Receive the Message

1. Create a JMS connection factory for the message producer using your broker's endpoint and then call the `createConnection` method against the factory.

```
// Create a connection factory.
final ActiveMQConnectionFactory connectionFactory = new
 ActiveMQConnectionFactory(wireLevelEndpoint);

// Pass the sign-in credentials.
connectionFactory.setUserName(activeMqUsername);
connectionFactory.setPassword(activeMqPassword);

// Establish a connection for the consumer.
final Connection consumerConnection = connectionFactory.createConnection();
consumerConnection.start();
```

> **Note**
>
> Message consumers should *never* use the `PooledConnectionFactory` class. For more information, see Always Use Connection Pooling.

2. Create a session, a queue named `MyQueue`, and a message consumer.

```
// Create a session.
final Session consumerSession = consumerConnection.createSession(false,
 Session.AUTO_ACKNOWLEDGE);

// Create a queue named "MyQueue".
final Destination consumerDestination = consumerSession.createQueue("MyQueue");

// Create a message consumer from the session to the queue.
final MessageConsumer consumer =
 consumerSession.createConsumer(consumerDestination);
```

3. Begin to wait for messages and receive the message when it arrives.

```
// Begin to wait for messages.
final Message consumerMessage = consumer.receive(1000);

// Receive the message when it arrives.
```

```
final TextMessage consumerTextMessage = (TextMessage) consumerMessage;
System.out.println("Message received: " + consumerTextMessage.getText());
```

> **Note**
>
> Unlike AWS messaging services (such as Amazon SQS), the consumer is constantly connected to the broker.

4. Close the consumer, session, and connection.

```
consumer.close();
consumerSession.close();
consumerConnection.close();
```

Integrating ActiveMQ brokers with LDAP

> **Important**
>
> LDAP integration is not supported for RabbitMQ brokers.

You can access your ActiveMQ brokers using the following protocols with TLS enabled:

- AMQP
- MQTT
- MQTT over WebSocket
- OpenWire
- STOMP
- STOMP over WebSocket

Amazon MQ offers a choice between native ActiveMQ authentication and LDAP authentication and authorization to manage user permissions. For information about restrictions related to ActiveMQ usernames and passwords, see Users.

To authorize ActiveMQ users and groups to works with queues and topics, you must edit your broker's configuration. Amazon MQ uses ActiveMQ's Simple Authentication Plugin to restrict

reading and writing to destinations. For more information and examples, see Always configure an authorization map and `authorizationEntry`.

> **Note**
>
> Currently, Amazon MQ doesn't support Client Certificate Authentication.

Topics

Integrate LDAP with ActiveMQ

You can authenticate Amazon MQ users through the credentials stored in your lightweight directory access protocol (LDAP) server. You can also add, delete, and modify Amazon MQ users and assign permissions to topics and queues through it. Management operations like creating, updating and deleting brokers still require IAM credentials and are not integrated with LDAP.

Customers who want to simplify and centralize their Amazon MQ broker authentication and authorization using an LDAP server can use this feature. Keeping all user credentials in the LDAP server saves time and effort by providing a central location for storing and managing these credentials.

Amazon MQ provides LDAP support using the Apache ActiveMQ JAAS plugin. Any LDAP server, such as Microsoft Active Directory or OpenLDAP that is supported by the plugin is also supported by Amazon MQ. For more information about the plugin, see the Security section of the Active MQ documentation.

In addition to users, you can specify access to topics and queues for a specific group or a user through your LDAP server. You do this by creating entries that represent topics and queues in your LDAP server and then assigning permissions to a specific LDAP user or a group. You can then configure broker to retrieve authorization data from the LDAP server.

Prerequisites

Before you add LDAP support to a new or existing Amazon MQ broker, you must set up a service account. This service account is required to initiate a connection to an LDAP server and must have the correct permissions to make this connection. This service account will set up LDAP authentication for your broker. Any successive client connections will be authenticated through the same connection.

The service account is an account in your LDAP server, which has access to initiate a connection. It is a standard LDAP requirement and you have to provide the service account credentials only once. After the connection is setup, all the future client connections are authenticated through your LDAP server. Your service account credentials are stored securely in an encrypted form, which is accessible only to Amazon MQ.

To integrate with ActiveMQ, a specific Directory Information Tree (DIT) is required on the LDAP server. For an example `ldif` file that clearly shows this structure, see *Import the following LDIF file into the LDAP server* in the Security section of the ActiveMQ documentation.

Getting Started with LDAP

To get started, navigate to the Amazon MQ console and choose **LDAP authentication and authorization** when you create a new Amazon MQ or edit an existing broker instance.

Provide the following information about the service account:

- **Fully qualified domain name** The location of the LDAP server to which authentication and authorization requests are to be issued.

> **Note**
>
> The fully qualified domain name of the LDAP server you supply must not include the protocol or port number. Amazon MQ will prepend the fully qualified domain name with the protocol `ldaps`, and will append the port number 636.
> For example, if you provide the following fully qualified domain: `example.com`, Amazon MQ will access your LDAP server using the following URL: `ldaps://example.com:636`.
> For the broker host to be able to successfully communicate with the LDAP server, the fully qualified domain name must be publicly resolvable. To keep the LDAP server private and secure, restrict inbound traffic in the server's inbound rules to only allow traffic originated from within the broker's VPC.

- **Service account username** The distinguished name of the user that will be used to perform the initial bind to the LDAP server.
- **Service account password** The password of the user performing the initial bind.

The following image highlights where to supply these details.

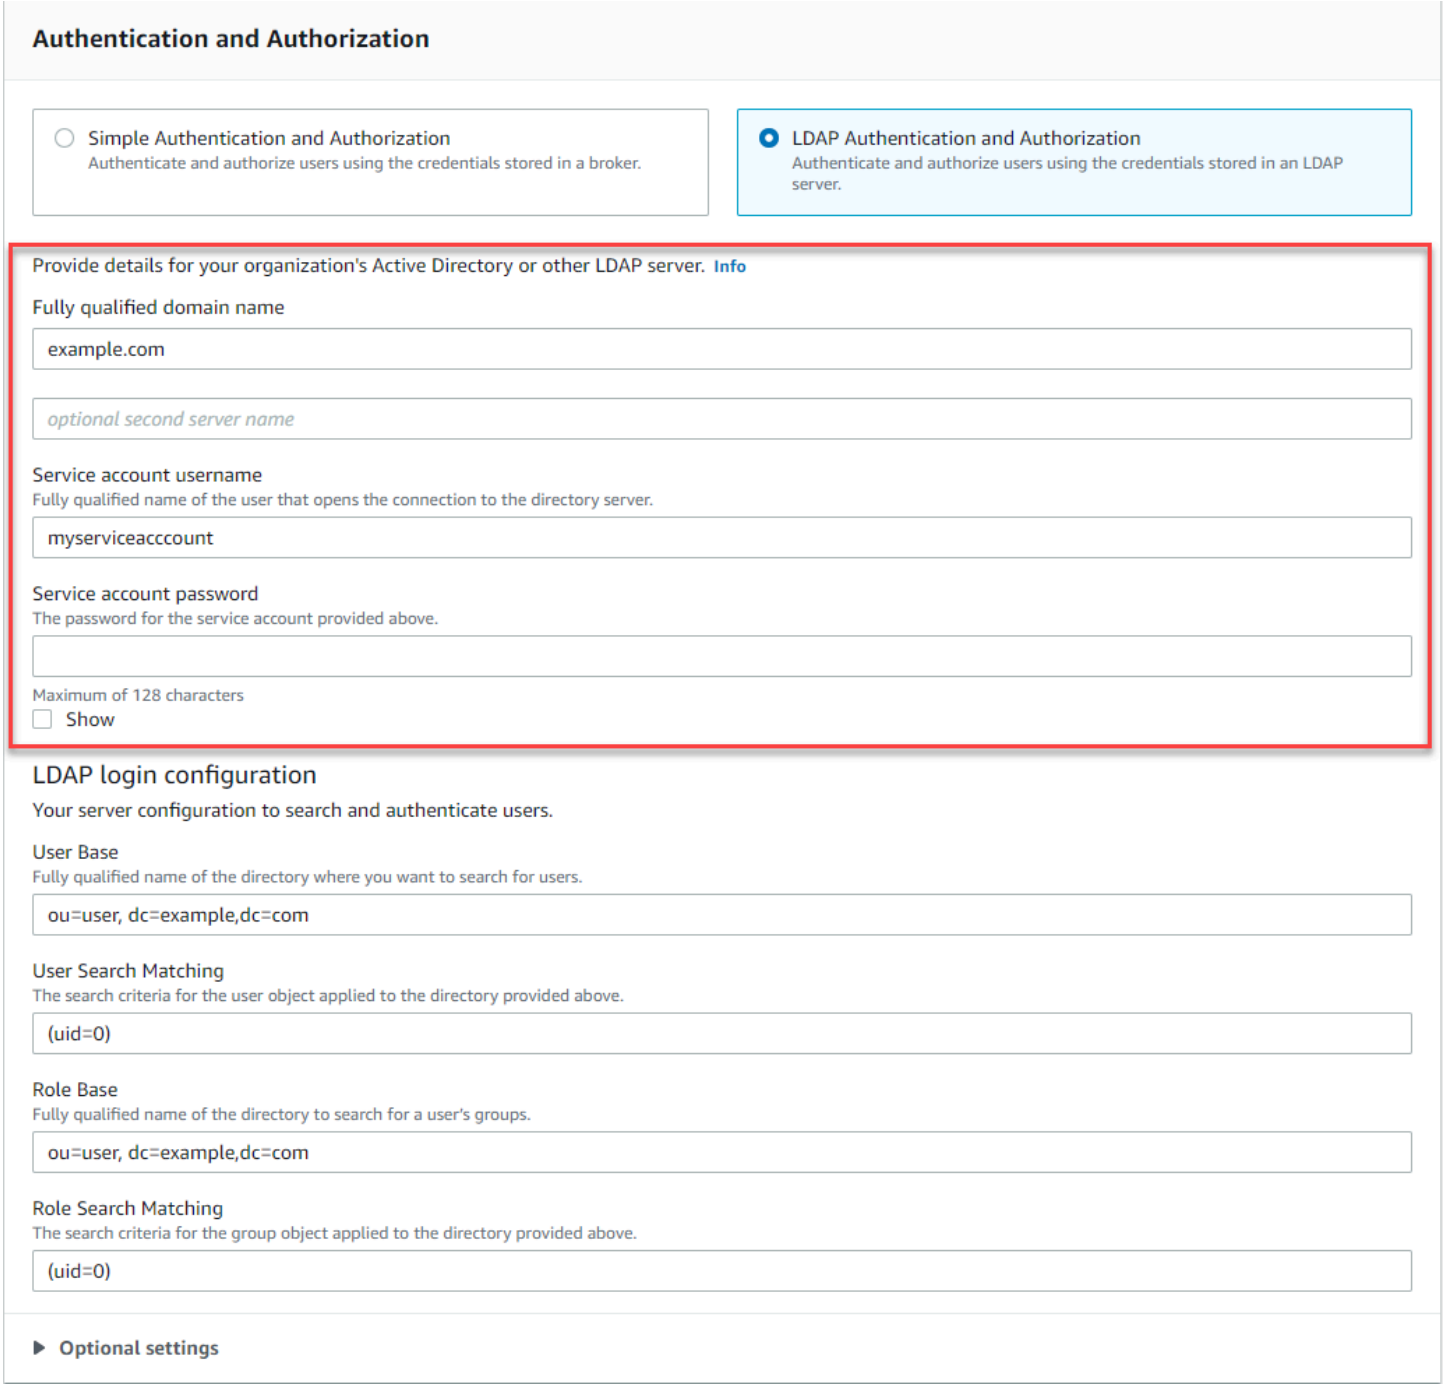

In the **LDAP login configuration** section, provide the following required information:

- **User Base** The distinguished name of the node in the directory information tree (DIT) that will be searched for users.
- **User Search Matching** The LDAP search filter that will be used to find users within the `userBase`. The client's username will be substituted into the `{0}` placeholder in the search filter. For more information, see Authentication and Authorization.
- **Role Base** The distinguished name of the node in the DIT that will be searched for roles. Roles can be configured as explicit LDAP group entries in your directory. A typical role entry may consist of one attribute for the name of the role, such as **common name (CN)**, and another attribute, such as `member`, with values representing the distinguished names or usernames of the users belonging to the role group. For example, given the organizational unit, `group`, you can provide the following distinguished name: `ou=group,dc=example,dc=com`.
- **Role Search Matching** The LDAP search filter that will be used to find roles within the `roleBase`. The distinguished name of the user matched by `userSearchMatching` will be substituted into the `{0}` placeholder in the search filter. The client's username will be substituted in place of the `{1}` placeholder. For example, if role entries in your directory include an attribute named `member`, containing the usernames for all users in that role, you can provide the following search filter: `(member:=uid={1})`.

The following image highlights where to specify these details.

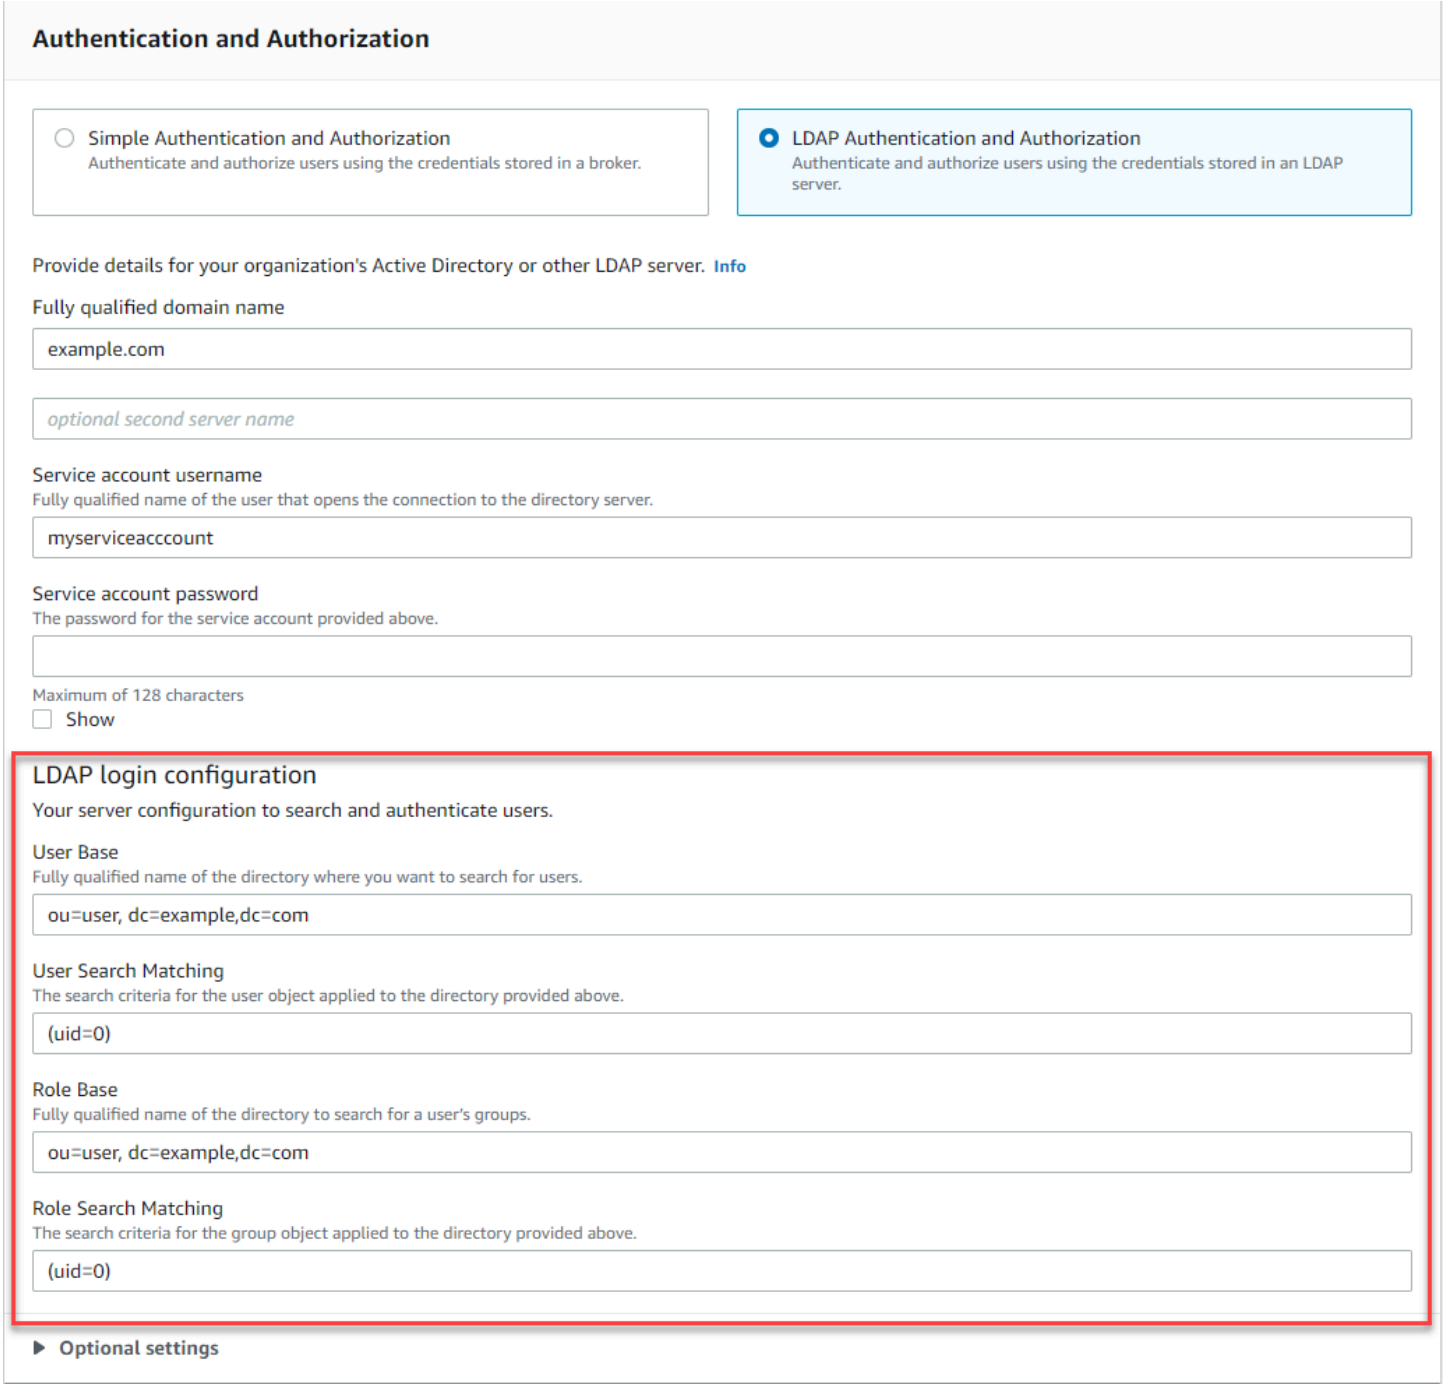

In the **Optional settings** section, you can provide the following optional information:

- **User Role Name** The name of the LDAP attribute in the user's directory entry for the user's group membership. In some cases, user roles may be identified by the value of an attribute in the user's directory entry. The `userRoleName` option allows you to provide the name of this attribute. For example, let's consider the following user entry:

```
dn: uid=jdoe,ou=user,dc=example,dc=com
objectClass: user
uid: jdoe
sn: jane
cn: Jane Doe
mail: j.doe@somecompany.com
memberOf: role1
userPassword: password
```

To provide the correct `userRoleName` for the example above, you would specify the `memberOf` attribute. If authentication is successful, the user is assigned the role `role1`.

- **Role Name** The group name attribute in a role entry whose value is the name of that role. For example, you can specify `cn` for a group entry's **common name**. If authentication succeeds, the user is assigned the the value of the `cn` attribute for each role entry that they are a member of.
- **User Search Subtree** Defines the scope for the LDAP user search query. If true, the scope is set to search the entire subtree under the node defined by `userBase`.
- **Role Search Subtree** Defines the scope for the LDAP role search query. If true, the scope is set to search the entire subtree under the node defined by `roleBase`.

The following image highlights where to specify these optional settings.

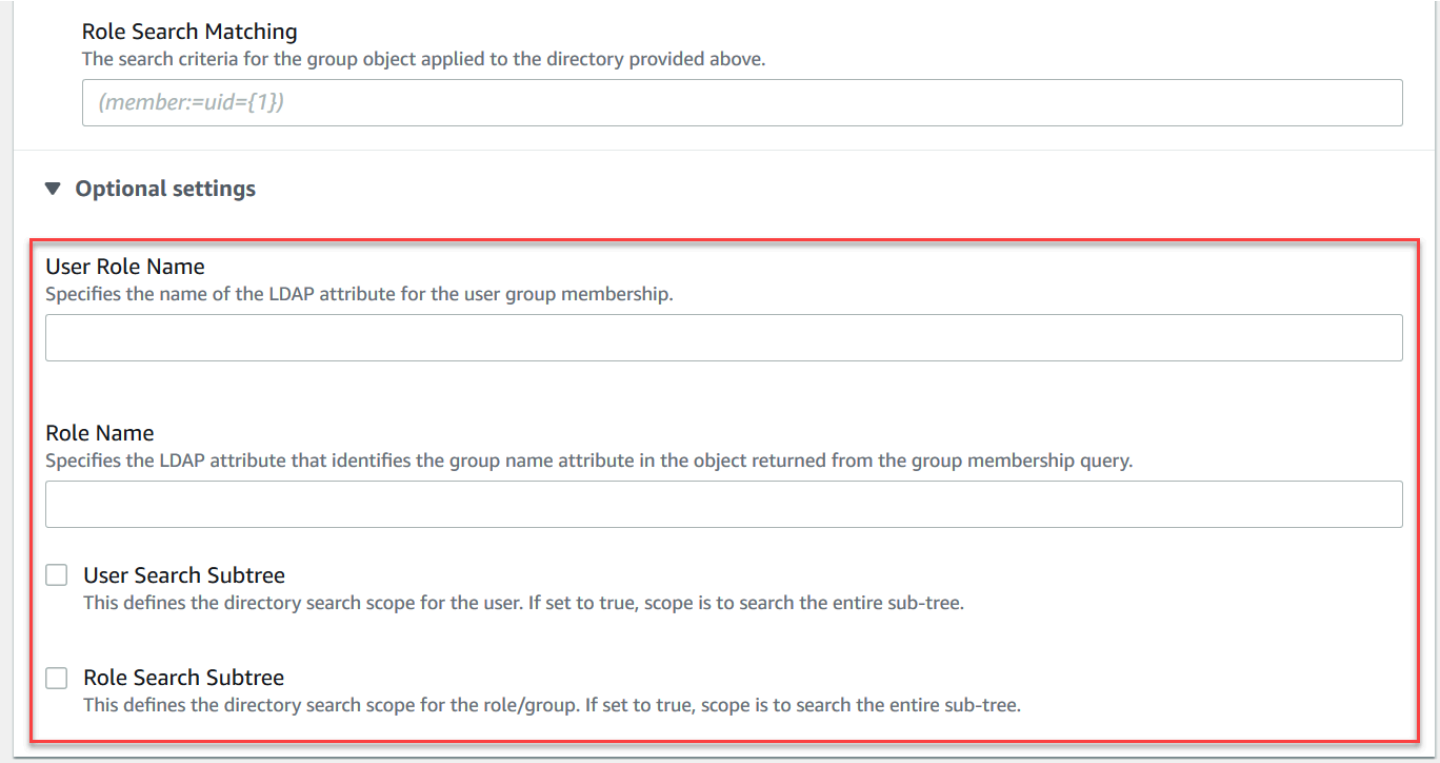

How LDAP integration works

You can think of integration in two main categories: the structure for authentication, and the structure for authorization.

Authentication

For authentication, client credentials must be valid. These credentials are validated against users in the user base in the LDAP server.

The user base supplied to the ActiveMQ broker must point to the node in the DIT where users are stored in the LDAP server. For example, if you are using AWS Managed Microsoft AD, and you have the domain components `corp`, `example`, and `com`, and within those you have organizational units `corp` and `Users`, you would use the following as your user base:

```
OU=Users,OU=corp,DC=corp,DC=example,DC=com
```

The ActiveMQ broker would search at this location in the DIT for users in order to authenticate client connection requests to the broker.

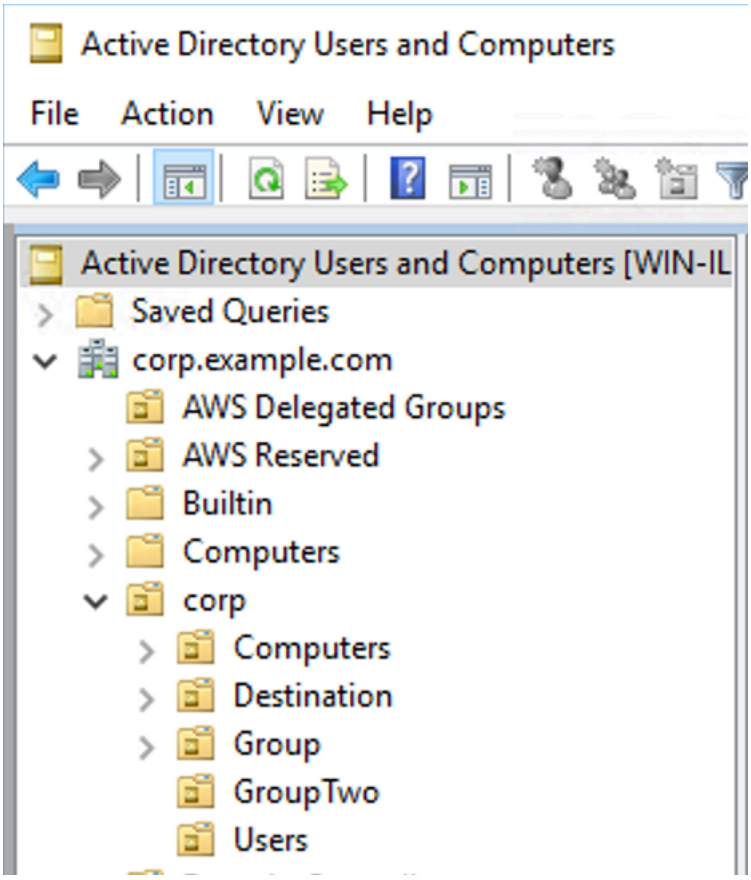

Because the ActiveMQ source code hardcodes the attribute name for users to `uid`, you must make sure that each user has this attribute set. For simplicity, you can use the user's connection username. For more information, see the activemq source code and Configuring ID mappings in Active Directory Users and Computers for Windows Server 2016 (and subsequent) versions.

To enable ActiveMQ console access for specific users, make sure they belong to the `amazonmq-console-admins` group.

Authorization

For authorization, permissions search bases are specified in the broker configuration. Authorization is done on a per-destination basis (or wildcard, destination set) via the `cachedLdapAuthorizationMap` element, found in the broker's `activemq.xml` configuration file. For more information, see Cached LDAP Authorization Module.

> **Note**
>
> To be able to use the `cachedLDAPAuthorizationMap` element in your broker's `activemq.xml` configuration file, you must choose the **LDAP Authentication and**

Authorization option when creating a configuration via the AWS Management Console, or set the `authenticationStrategy` property to LDAP when creating a new configuration using the Amazon MQ API.

You must provide the following three attributes as part of the `cachedLDAPAuthorizationMap` element:

- `queueSearchBase`
- `topicSearchBase`
- `tempSearchBase`

Important

To prevent sensitive information from being directly placed in the broker's configuration file, Amazon MQ blocks the following attributes from being used in `cachedLdapAuthorizationMap`:

- `connectionURL`
- `connectionUsername`
- `connectionPassword`

When you create a broker, Amazon MQ substitutes the values you provide via the AWS Management Console, or in the `ldapServerMetadata` property of your API request, for the above attributes.

The following demonstrates a working example of the `cachedLdapAuthorizationMap`.

```
<authorizationPlugin>
    <map>
        <cachedLDAPAuthorizationMap
            queueSearchBase="ou=Queue,ou=Destination,ou=corp,dc=corp,dc=example,dc=com"
            topicSearchBase="ou=Topic,ou=Destination,ou=corp,dc=corp,dc=example,dc=com"
            tempSearchBase="ou=Temp,ou=Destination,ou=corp,dc=corp,dc=example,dc=com"
            refreshInterval="300000"
            legacyGroupMapping="false"
```

```
        />
    </map>
</authorizationPlugin>
```

These values identify the locations within the DIT where permissions for each type of destination are specified. So for the above example with AWS Managed Microsoft AD, using the same domain components of `corp`, `example`, and `com`, you would specify an organizational unit named `destination` to contain all your destination types. Within that OU, you would create one for `queues`, one for `topics`, and one for `temp` destinations.

This would mean that your queue search base, which provides authorization information for destinations of type queue, would have the following location in your DIT:

```
OU=Queue,OU=Destination,OU=corp,DC=corp,DC=example,DC=com
```

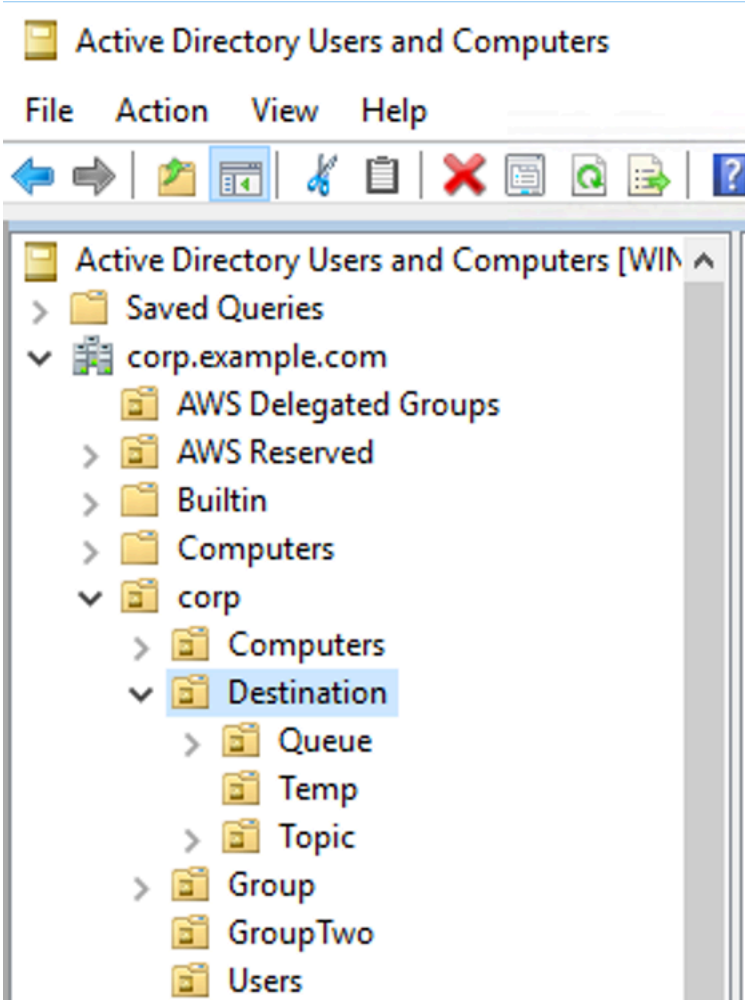

Similarly, permissions rules for topics and temp destinations would be located at the same level in the DIT:

```
OU=Topic,OU=Destination,OU=corp,DC=corp,DC=example,DC=com
OU=Temp,OU=Destination,OU=corp,DC=corp,DC=example,DC=com
```

Within the OU for each type of destination (queue, topic, temp), either a wildcard or specific destination name can be provided. For example, to provide an authorization rule for all queues that start with the prefix DEMO.EVENTS.$., you could create the following OU:

```
OU=DEMO.EVENTS.$,OU=Queue,OU=Destination,OU=corp,DC=corp,DC=example,DC=com
```

> **Note**
>
> The `DEMO.EVENTS.$` OU is within the `Queue` OU.

For more info on wildcards in ActiveMQ, see Wildcards

To provide authorization rules for specific queues, such as DEMO.MYQUEUE, specify something like the following:

```
OU=DEMO.MYQUEUE,OU=Queue,OU=Destination,OU=corp,DC=corp,DC=example,DC=com
```

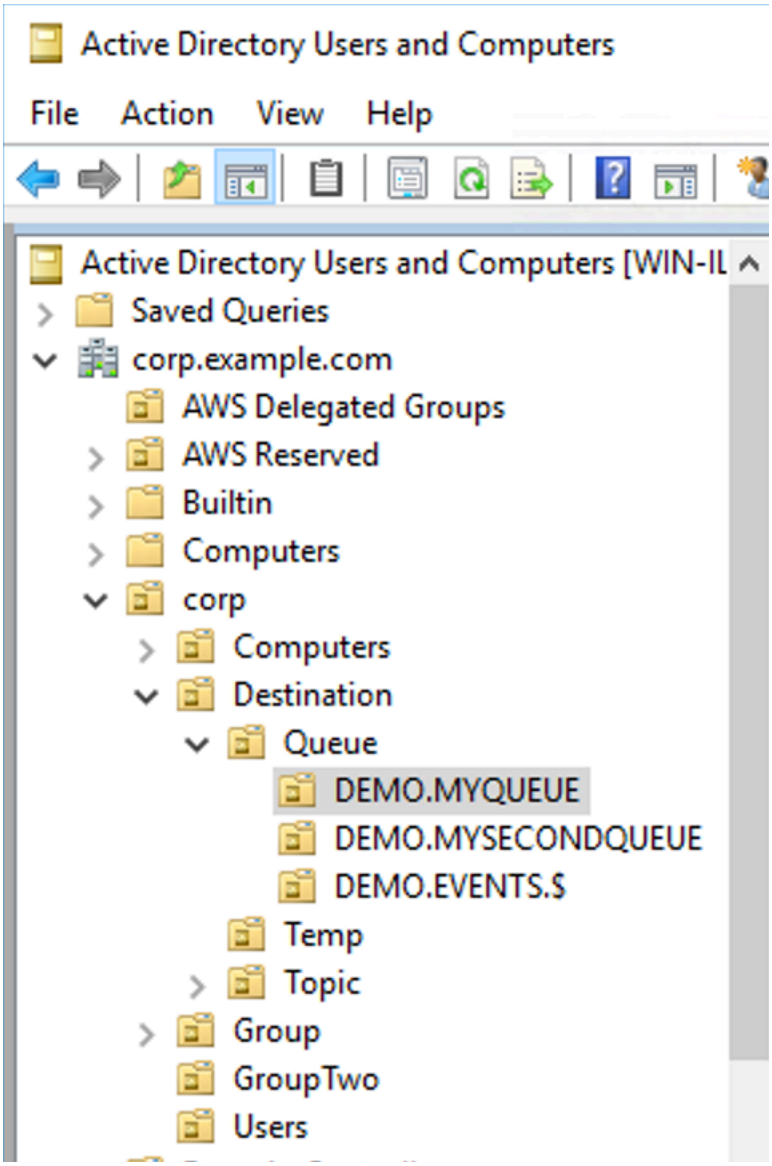

Security Groups

Within each OU that represents a destination or a wildcard, you must create three security groups. As with all permissions in ActiveMQ, these are read/write/admin permissions. For more information on what each of these permissions allows a user to do, see Security in the ActiveMQ documentation.

You must name these security groups `read`, `write`, and `admin`. Within each of these security groups, you can add users or groups, who will then have permission to perform the associated actions. You'll need these security groups for each wildcard destination set or individual destination.

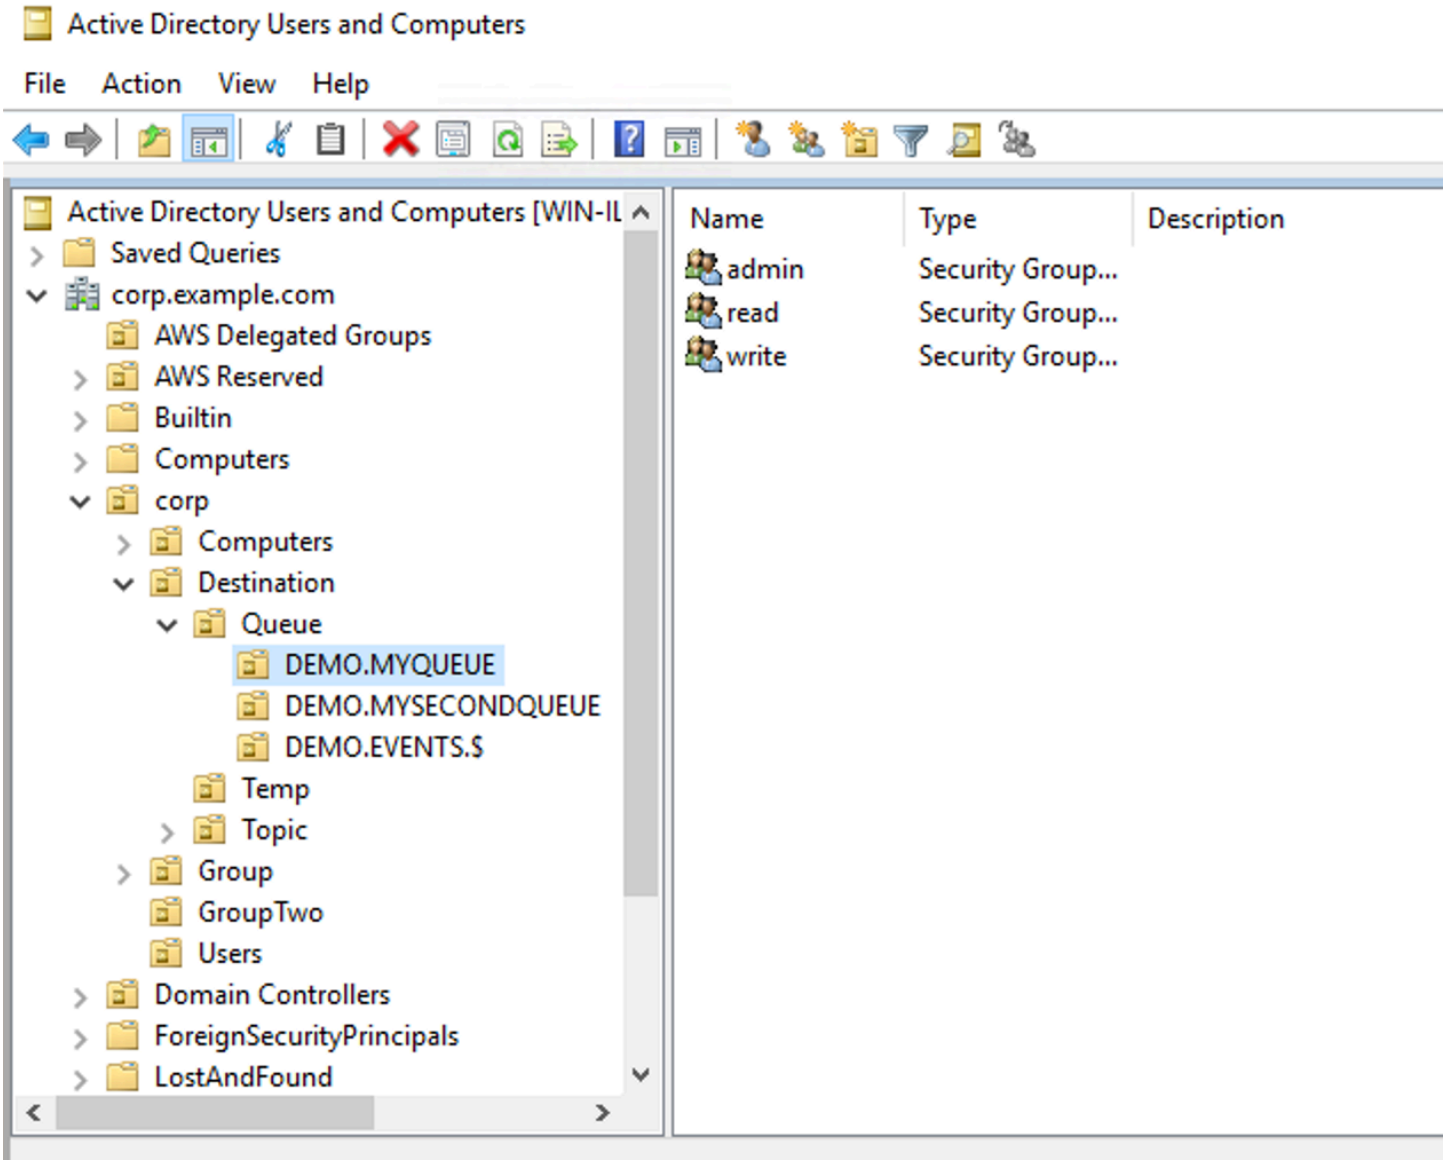

Note

When you create the admin group, a conflict will arise with the group name. This conflict happens because the legacy pre-Windows 2000 rules do not allow groups to share the same name, even if the groups are in different locations of the DIT. The value in the **pre-Windows 2000** text box has no impact on the setup, but it must be globally unique. To avoid this conflict, you can append a `uuid` suffix to each `admin` group.

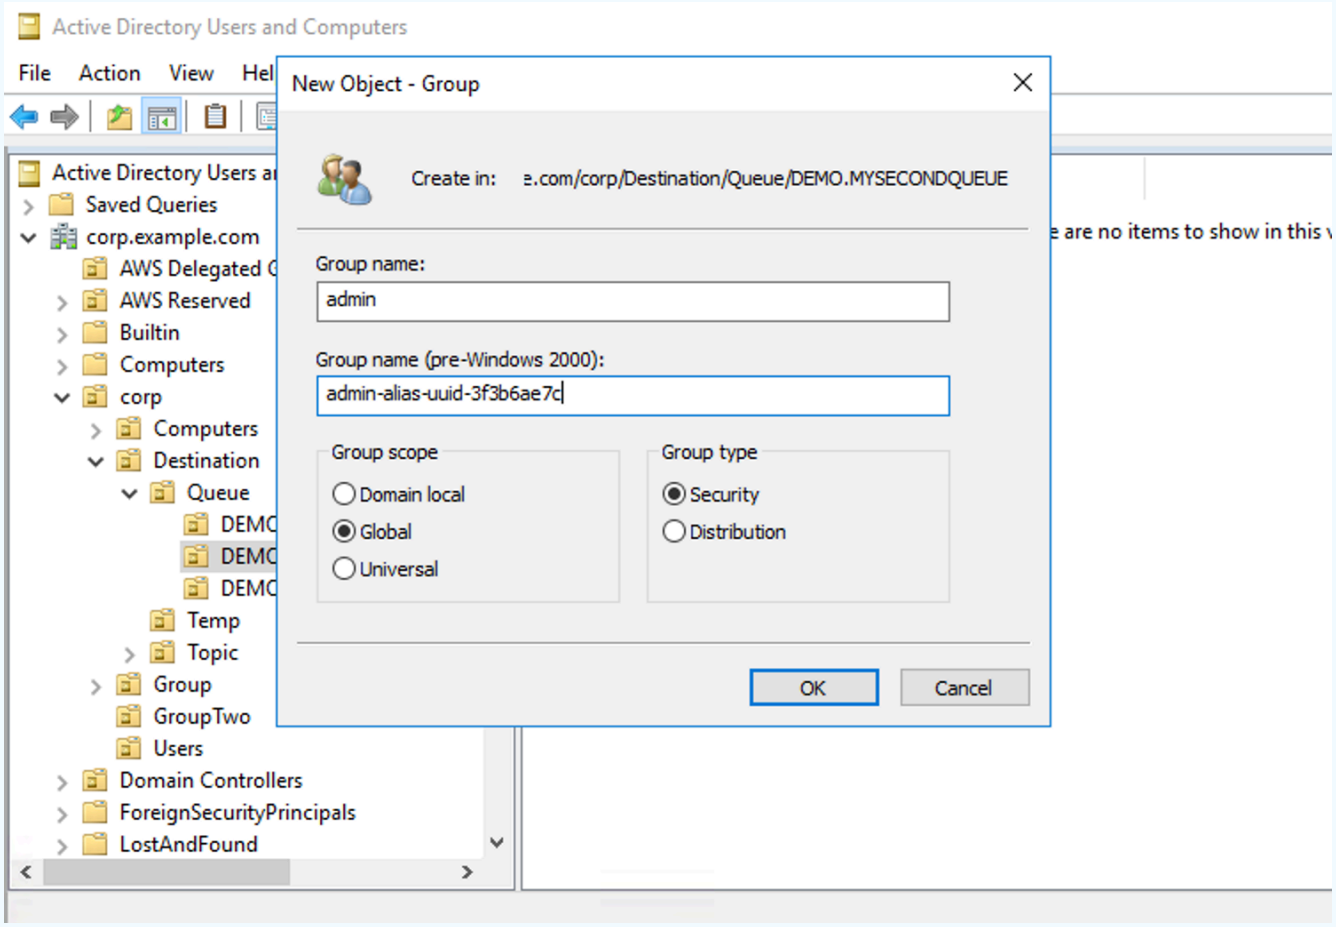

Adding a user to the `admin` security group for a particular destination will enable the user to create and delete that topic. Adding them to the `read` security group will enable them to read from the destination, and adding them to the `write` group will enable them to write to the destination.

In addition to adding individual users to security group permissions, you can also add entire groups. However, because ActiveMQ again hardcodes attribute names for groups, you must ensure the group you want to add has the object class `groupOfNames`, as shown in the activemq source code.

To do this, follow the same process as with the `uid` for users. See Configuring ID mappings in Active Directory Users and Computers for Windows Server 2016 (and subsequent) versions.

Creating and managing ActiveMQ broker users

An ActiveMQ *user* is a person or an application that can access the queues and topics of an ActiveMQ broker. You can configure users to have specific permissions. For example, you can allow some users to access the ActiveMQ Web Console.

A *group* is a semantic label. You can assign a group to a user and configure permissions for groups to send to, receive from, and administer specific queues and topics.

> **Note**
>
> You can't configure groups independently of users. A group label is created when you add at least one user to it and deleted when you remove all users from it.

The following examples show how you can create, edit, and delete Amazon MQ broker users using the AWS Management Console.

Topics

- To create a new user
- To edit an existing user
- To delete an existing user

To create a new user

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**) and then choose **View details**.

   On the ***MyBroker*** page, in the **Users** section, all the users for this broker are listed.

	Username	Console access	Groups	Pending modifications
○	paolo.santos	No	Devs	
○	jane.doe	Yes	Admins	

3. Choose **Create user**.
4. In the **Create user** dialog box, type a **Username** and **Password**.
5. (Optional) Type the names of groups to which the user belongs, separated by commas (for example: `Devs, Admins`).
6. (Optional) To enable the user to access the ActiveMQ Web Console, choose **ActiveMQ Web Console**.
7. Choose **Create user**.

> **Important**
>
> Making changes to a user does *not* apply the changes to the user immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.

To edit an existing user

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**) and then choose **View details**.

   On the ***MyBroker*** page, in the **Users** section, all the users for this broker are listed.

	Username	Console access	Groups	Pending modifications
○	paolo.santos	No	Devs	
○	jane.doe	Yes	Admins	

3. Select your sign-in credentials and choose **Edit**.

   The **Edit user** dialog box is displayed.
4. (Optional) Type a new **Password**.
5. (Optional) Add or remove the names of groups to which the user belongs, separated by commas (for example: `Managers, Admins`).
6. (Optional) To enable the user to access the ActiveMQ Web Console, choose **ActiveMQ Web Console**.
7. To save the changes to the user, choose **Done**.

> **Important**
>
> Making changes to a user does *not* apply the changes to the user immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.

To delete an existing user

1. Sign in to the Amazon MQ console.
2. From the broker list, choose the name of your broker (for example, **MyBroker**) and then choose **View details**.

   On the ***MyBroker*** page, in the **Users** section, all the users for this broker are listed.

	Username	Console access	Groups	Pending modifications
○	paolo.santos	No	Devs	
○	jane.doe	Yes	Admins	

3. Select a your sign-in credentials (for example, ***MyUser***) and then choose **Delete**.
4. To confirm deleting the user, in the **Delete *MyUser*?** dialog box, choose **Delete**.

> **Important**
>
> Making changes to a user does *not* apply the changes to the user immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.

Amazon MQ for ActiveMQ best practices

Use this as a reference to quickly find recommendations for maximizing performance and minimizing throughput costs when working with ActiveMQ brokers on Amazon MQ.

Topics

- Connecting to Amazon MQ

- Ensuring effective Amazon MQ performance
- Avoid slow restarts by recovering prepared XA transactions

Connecting to Amazon MQ

The following design patterns can improve the effectiveness of your application's connection to your Amazon MQ broker.

Topics

Never Modify or Delete the Amazon MQ Elastic Network Interface

When you first create an Amazon MQ broker, Amazon MQ provisions an elastic network interface in the Virtual Private Cloud (VPC) under your account and, thus, requires a number of EC2 permissions. The network interface allows your client (producer or consumer) to communicate with the Amazon MQ broker. The network interface is considered to be within the *service scope* of Amazon MQ, despite being part of your account's VPC.

> **Warning**
>
> You must not modify or delete this network interface. Modifying or deleting the network interface can cause a permanent loss of connection between your VPC and your broker.

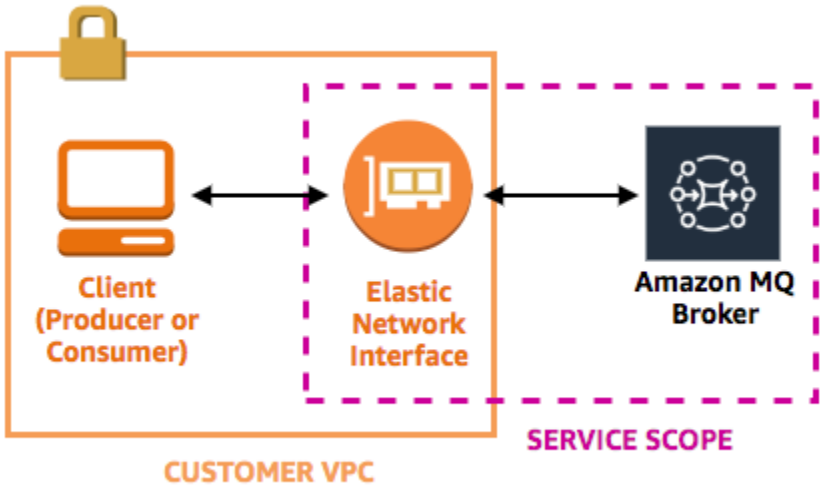

Always Use Connection Pooling

In a scenario with a single producer and single consumer (such as the Getting Started with Amazon MQ tutorial), you can use a single `ActiveMQConnectionFactory` class for every producer and consumer. For example:

```
// Create a connection factory.
final ActiveMQConnectionFactory connectionFactory = new
 ActiveMQConnectionFactory(wireLevelEndpoint);

// Pass the sign-in credentials.
connectionFactory.setUserName(activeMqUsername);
connectionFactory.setPassword(activeMqPassword);

// Establish a connection for the consumer.
final Connection consumerConnection = connectionFactory.createConnection();
consumerConnection.start();
```

However, in more realistic scenarios with multiple producers and consumers, it can be costly and inefficient to create a large number of connections for multiple producers. In these scenarios, you should group multiple producer requests using the `PooledConnectionFactory` class. For example:

> **Note**
>
> Message consumers should *never* use the `PooledConnectionFactory` class.

```
// Create a connection factory.
final ActiveMQConnectionFactory connectionFactory = new
 ActiveMQConnectionFactory(wireLevelEndpoint);

// Pass the sign-in credentials.
connectionFactory.setUserName(activeMqUsername);
connectionFactory.setPassword(activeMqPassword);

// Create a pooled connection factory.
final PooledConnectionFactory pooledConnectionFactory = new PooledConnectionFactory();
pooledConnectionFactory.setConnectionFactory(connectionFactory);
pooledConnectionFactory.setMaxConnections(10);

// Establish a connection for the producer.
final Connection producerConnection = pooledConnectionFactory.createConnection();
producerConnection.start();
```

Always Use the Failover Transport to Connect to Multiple Broker Endpoints

If you need your application to connect to multiple broker endpoints—for example, when you use an active/standby deployment mode or when you migrate from an on-premises message broker to Amazon MQ—use the Failover Transport to allow your consumers to randomly connect to either one. For example:

```
failover:(ssl://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
east-2.amazonaws.com:61617,ssl://b-9876l5k4-32ji-109h-8gfe-7d65c4b132a1-2.mq.us-
east-2.amazonaws.com:61617)?randomize=true
```

Avoid Using Message Selectors

It is possible to use JMS selectors to attach filters to topic subscriptions (to route messages to consumers based on their content). However, the use of JMS selectors fills up the Amazon MQ broker's filter buffer, preventing it from filtering messages.

In general, avoid letting consumers route messages because, for optimal decoupling of consumers and producers, both the consumer and the producer should be ephemeral.

Prefer Virtual Destinations to Durable Subscriptions

A durable subscription can help ensure that the consumer receives all messages published to a topic, for example, after a lost connection is restored. However, the use of durable subscriptions also precludes the use of competing consumers and might have performance issues at scale. Consider using virtual destinations instead.

If using Amazon VPC peering, avoid client IPs in CIDR range `10.0.0.0/16`

If you are setting up Amazon VPC peering between on-premise infrastructure and your Amazon MQ broker, you must not configure client connections with IPs in CIDR range `10.0.0.0/16`.

Ensuring effective Amazon MQ performance

The following design patterns can improve the effectiveness and performance of your Amazon MQ broker.

Topics

Disable Concurrent Store and Dispatch for Queues with Slow Consumers

By default, Amazon MQ optimizes for queues with fast consumers:

- Consumers are considered *fast* if they are able to keep up with the rate of messages generated by producers.
- Consumers are considered *slow* if a queue builds up a backlog of unacknowledged messages, potentially causing a decrease in producer throughput.

To instruct Amazon MQ to optimize for queues with slow consumers, set the `concurrentStoreAndDispatchQueues` attribute to `false`. For an example configuration, see `concurrentStoreAndDispatchQueues`.

Choose the Correct Broker Instance Type for the Best Throughput

The message throughput of a broker instance type depends on your application's use case and the following factors:

- Use of ActiveMQ in persistent mode
- Message size
- The number of producers and consumers
- The number of destinations

Understanding the relationship between message size, latency, and throughput

Depending on your use case, a larger broker instance type might not necessarily improve system throughput. When ActiveMQ writes messages to durable storage, the size of your messages determines your system's limiting factor:

- If your messages are smaller than 100 KB, persistent storage latency is the limiting factor.
- If your messages are larger than 100 KB, persistent storage throughput is the limiting factor.

When you use ActiveMQ in persistent mode, writing to storage normally occurs when there are either few consumers or when the consumers are slow. In non-persistent mode, writing to storage also occurs with slow consumers if the heap memory of the broker instance is full.

To determine the best broker instance type for your application, we recommend testing different broker instance types. For more information, see Broker instance types and also Measuring the Throughput for Amazon MQ using the JMS Benchmark.

Use cases for larger broker instance types

There are three common use cases when larger broker instance types improve throughput:

- **Non-persistent mode** – When your application is less sensitive to losing messages during broker instance failover (for example, when broadcasting sports scores), you can often use ActiveMQ's non-persistent mode. In this mode, ActiveMQ writes messages to persistent storage only if the heap memory of the broker instance is full. Systems that use non-persistent mode can benefit from the higher amount of memory, faster CPU, and faster network available on larger broker instance types.

- **Fast consumers** – When active consumers are available and the `concurrentStoreAndDispatchQueues` flag is enabled, ActiveMQ allows messages to flow directly from producer to consumer without sending messages to storage (even in persistent mode). If your application can consume messages quickly (or if you can design your consumers to do this), your application can benefit from a larger broker instance type. To let your application consume messages more quickly, add consumer threads to your application instances or scale up your application instances vertically or horizontally.
- **Batched transactions** – When you use persistent mode and send multiple messages per transaction, you can achieve an overall higher message throughput by using larger broker instance types. For more information, see Should I Use Transactions? in the ActiveMQ documentation.

Choose the correct broker storage type for the best throughput

To take advantage of high durability and replication across multiple Availability Zones, use Amazon EFS. To take advantage of low latency and high throughput, use Amazon EBS. For more information, see Storage.

Configure Your Network of Brokers Correctly

When you create a network of brokers, configure it correctly for your application:

- **Enable persistent mode** – Because (relative to its peers) each broker instance acts like a producer or a consumer, networks of brokers don't provide distributed replication of messages. The first broker that acts as a consumer receives a message and persists it to storage. This broker sends an acknowledgement to the producer and forwards the message to the next broker. When the second broker acknowledges the persistence of the message, the first broker deletes the message.

  If persistent mode is disabled, the first broker acknowledges the producer without persisting the message to storage. For more information, see Replicated Message Store and What is the difference between persistent and non-persistent delivery? in the Apache ActiveMQ documentation.
- **Don't disable advisory messages for broker instances** – For more information, see Advisory Message in the Apache ActiveMQ documentation.

- **Don't use multicast broker discovery** – Amazon MQ doesn't support broker discovery using multicast. For more information, see What is the difference between discovery, multicast, and zeroconf? in the Apache ActiveMQ documentation.

Avoid slow restarts by recovering prepared XA transactions

ActiveMQ supports distributed (XA) transactions. Knowing how ActiveMQ processes XA transactions can help avoid slow recovery times for broker restarts and failovers in Amazon MQ

Unresolved prepared XA transactions are replayed on every restart. If these remain unresolved, their number will grow over time, significantly increasing the time needed to start up the broker. This affects restart and failover time. You must resolve these transactions with a `commit()` or a `rollback()` so that performance doesn't degrade over time.

To monitor your unresolved prepared XA transactions, you can use the `JournalFilesForFastRecovery` metric in Amazon CloudWatch Logs. If this number is increasing, or is consistently higher than 1, you should recover your unresolved transactions with code similar to the following example. For more information, see Quotas in Amazon MQ.

The following example code walks through prepared XA transactions and closes them with a `rollback()`.

```
import org.apache.activemq.ActiveMQXAConnectionFactory;

import javax.jms.XAConnection;
import javax.jms.XASession;
import javax.transaction.xa.XAResource;
import javax.transaction.xa.Xid;

public class RecoverXaTransactions {
    private static final ActiveMQXAConnectionFactory ACTIVE_MQ_CONNECTION_FACTORY;
    final static String WIRE_LEVEL_ENDPOINT =
            "tcp://localhost:61616";;
    static {
        final String activeMqUsername = "MyUsername123";
        final String activeMqPassword = "MyPassword456";
        ACTIVE_MQ_CONNECTION_FACTORY = new
 ActiveMQXAConnectionFactory(activeMqUsername, activeMqPassword, WIRE_LEVEL_ENDPOINT);
        ACTIVE_MQ_CONNECTION_FACTORY.setUserName(activeMqUsername);
        ACTIVE_MQ_CONNECTION_FACTORY.setPassword(activeMqPassword);
```

```
    }

    public static void main(String[] args) {
        try {
            final XAConnection connection =
 ACTIVE_MQ_CONNECTION_FACTORY.createXAConnection();
            XASession xaSession = connection.createXASession();
            XAResource xaRes = xaSession.getXAResource();

            for (Xid id : xaRes.recover(XAResource.TMENDRSCAN)) {
                xaRes.rollback(id);
            }
            connection.close();

        } catch (Exception e) {
        }
    }
}
```

In a real-world scenario, you could check your prepared XA transactions against your XA Transaction Manager. Then you can decide whether to handle each prepared transaction with a `rollback()` or a `commit()`.

Cross-Region data replication for Amazon MQ for ActiveMQ

Amazon MQ for ActiveMQ offers a Cross-Region data replication (CRDR) feature that allows for asynchronous message replication from the primary broker in a primary AWS Region to the replica broker in a replica Region. By issuing a failover request to the Amazon MQ API, the current replica broker is promoted to the primary broker role, and the current primary broker is demoted to the replica role.

This section provides tutorials on how to set up Cross-Region data replication with Amazon MQ for ActiveMQ.

Topics

Primary and replica brokers in Amazon MQ

You can create primary and replica brokers for asynchronous data replication from the primary broker in a primary AWS Region to the replica broker in a replica Region. The *primary Region* consists of a redundant pair of active/standby brokers referred to as the *primary broker*. The *secondary Region* consists of a redundant pair of active/standby brokers referred to as the *replica broker*.

The following diagram illustrates a replica broker in a secondary Region receiving asynchronous replicated data from the primary broker in the primary Region.

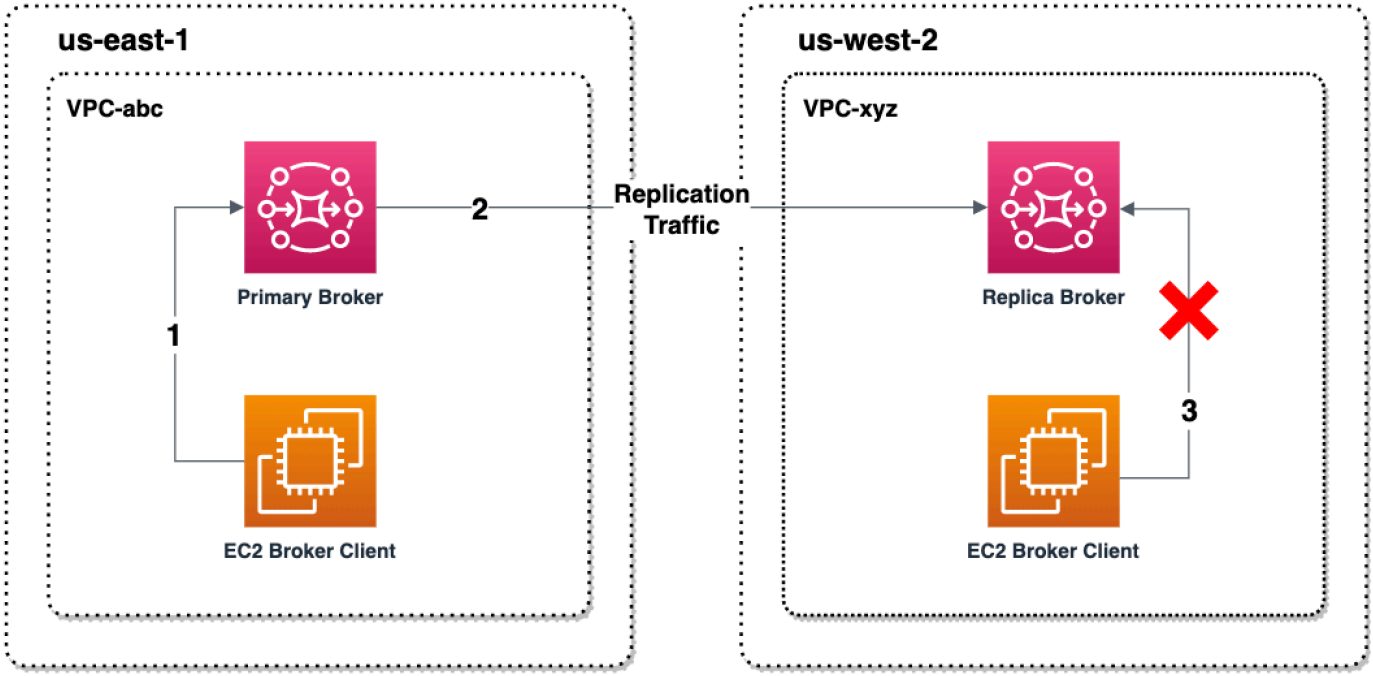

Primary and replica brokers act as a cross-Region data recovery solution. If the primary broker in the primary Region fails, you can promote the replica broker in the secondary Region to primary by initiating a switchover or failover. The former primary broker then becomes the replica broker, and the former replica broker is promoted to primary broker. For instructions on creating a primary and replica broker, see Creating and deleting a Cross-Region data replication broker.

> **Note**
>
> Only available for active/standby brokers.

Creating and deleting a Cross-Region data replication broker

With Cross-Region data replication (CRDR), you can switch between Amazon MQ for ActiveMQ message brokers in two AWS Regions as needed. You can designate an existing broker as a primary broker and create a replica for this broker, or create a new primary and replica broker together. You can then promote the replica broker to the primary broker role using the Amazon MQ `Promote` API operation. For more information about primary and replica brokers, see Primary and replica brokers in Amazon MQ.

The following instructions describe how you can create and configure a replica broker using the Amazon MQ Management Console.

Topics

Prerequisites

To use the cross-Region data replication feature, you must review and comply with the following prerequisites:

- **Version**: The cross-Region data replication feature is only available for Amazon MQ for ActiveMQ brokers on versions 5.17.6 and above.
- **Region**: Cross-Region data replication is supported in the following regions: US East (Ohio), US East (N. Virginia), US West (Oregon), and US West (N. California).
- **Instance type**: Cross-Region data replication is only available for broker instance sizes `mq.m5.large` and above.
- **Deployment type**: Cross-Region data replication is only available for active/standby brokers with multi-availability zone deployment.
- **Broker status**: You can only create a replica broker for a primary broker with the broker status `Running`.

Step 1 (Optional): Create a new primary broker

Create a new primary broker

1. Sign in to the Amazon MQ console.
2. On the Brokers page of the Amazon MQ console, choose **Create brokers**.
3. On the **Select broker engine** page, choose **Apache ActiveMQ**.
4. On the **Select deployment and storage** page, in the **Deployment mode and storage type** section, do the following:
   - For the **Deployment mode**, choose **Active/standby broker**. An **Active/standby broker** is comprised of two brokers in two different Availability Zones configured in a redundant pair. These brokers communicate synchronously with your application and with Amazon EFS. For more information, see Broker Architecture.
5. Choose **Next**.
6. On the **Configure settings** page, in the **Details** section, do the following:
   a. Enter the **Broker name**.

      > **Important**
      > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker names. Broker names are accessible to other AWS services, including CloudWatch Logs. Broker names are not intended to be used for private or sensitive data.

   b. Choose the **Broker instance type** (for example, **mq.m5.large**). For more information, see Broker instance types.
7. In the **ActiveMQ Web Console access** section, provide a **Username** and **Password**. The following restrictions apply to broker usernames and passwords:
   - Your username can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~).
   - Your password must be at least 12 characters long, contain at least 4 unique characters and must not contain commas, colons, or equal signs (,:=).

> **Important**
>
> Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

The green flash bar at the top of the page confirms that Amazon MQ is creating the replica broker in the recovery Region. You can also see the CRDR role and RPO status for your brokers. To turn off the CRDR Role and RPO Status columns, choose the gear icon in the top right corner of the **Brokers** table. Then, on the **Preferences** page, turn off CRDR Role or RPO Status.

Step 2: Create a replica of an existing broker

1. On the Brokers page of the Amazon MQ console, choose **Create replica broker**.
2. On the **Choose primary broker page**, select an existing broker to use as a CRDR primary broker. Then, choose **Next**.
3. On the **Configure replica broker** page, use the drop down menu to choose the replica Region.
4. In the **ActiveMQ console user for replica broker** section, provide a **Username** and **Password** for the replica broker console user. The following restrictions apply to broker usernames and passwords:

   - Your username can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~).
   - Your password must be at least 12 characters long, contain at least 4 unique characters and must not contain commas, colons, or equal signs (,:=).

> **Important**
>
> Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

5. In the **Data replication user to bridge access between brokers** section, provide a **Username** and **Password** for the user that will access both the primary and replica broker. The following restrictions apply to broker usernames and passwords:

   - Your username can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~).
   - Your password must be at least 12 characters long, contain at least 4 unique characters and must not contain commas, colons, or equal signs (,:=).

   > **Important**
   >
   > Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

   Configure any additional settings. Then, choose **Next**.
6. On the **Review and create** page, review the replica broker details. Then, choose **Create replica broker**.
7. Next, reboot the primary broker. This will also reboot the replica broker. For instructions on rebooting your broker, see Rebooting a Broker.

For more information on configuring additional settings for your ActiveMQ broker, see Creating and connecting to an ActiveMQ broker

Delete a CRDR broker

To delete a primary or replica CRDR broker, you must first unpair then reboot the brokers. The following instructions show how you can unpair and reboot the brokers using the AWS Management Console.

1. On the **Brokers** page, select the CRDR broker you want to unpair, then choose **Edit**.
2. On the broker **Edit** page in the **Data replication** section, choose **Unpair brokers**.
3. Enter "unpair" in the pop-up window to confirm your choice. Then choose **Unpair brokers**.

4. Next, reboot the unpaired primary broker. This will also reboot the replica broker. For instructions on rebooting your broker, see Rebooting a Broker. After the primary broker is rebooted, both brokers are unpaired and can be individually deleted. To delete your broker, see Deleting a broker.

Initiating switchover or failover to promote replica broker to primary broker role

You can initiate a switchover or failover when you want to promote the replica broker to the primary broker role. When you promote the replica broker, the primary broker is demoted to the replica broker role.

A **switchover** prioritizes consistency over availability. Brokers are guaranteed to have identical state when this failover operation completes. With a switchover, there may be a period where neither broker is available for client connections while inter-broker consistency is established. Both brokers will have the same state at the instant when the replica is promoted. Switchover success depends on the health of both regions and the inter-region network to succeed.

A **failover** prioritizes availability over consistency. Brokers are not guaranteed to have identical states when this operation completes. With a failover, the replica broker is guaranteed to become immediately available to serve client traffic, without waiting for any replication data to be synchronized, or for the primary to receive the shutdown signal. Failover depends on neither the health of the original primary region nor the inter-region network to succeed.

The following diagram illustrates a switchover in which neither broker accepts client connections while the replication queue is being drained and broker states are synchronized. In this process, the client in the primary broker's VPC is unable to produce further state changes while the operation is in progress, and the primary broker is being demoted to a replica. When the replication queue is drained and the two brokers achieve identical state, the client in the replica broker's VPC is unable to connect to the replica broker until the failover operation completes, and the replica broker is promoted to primary.

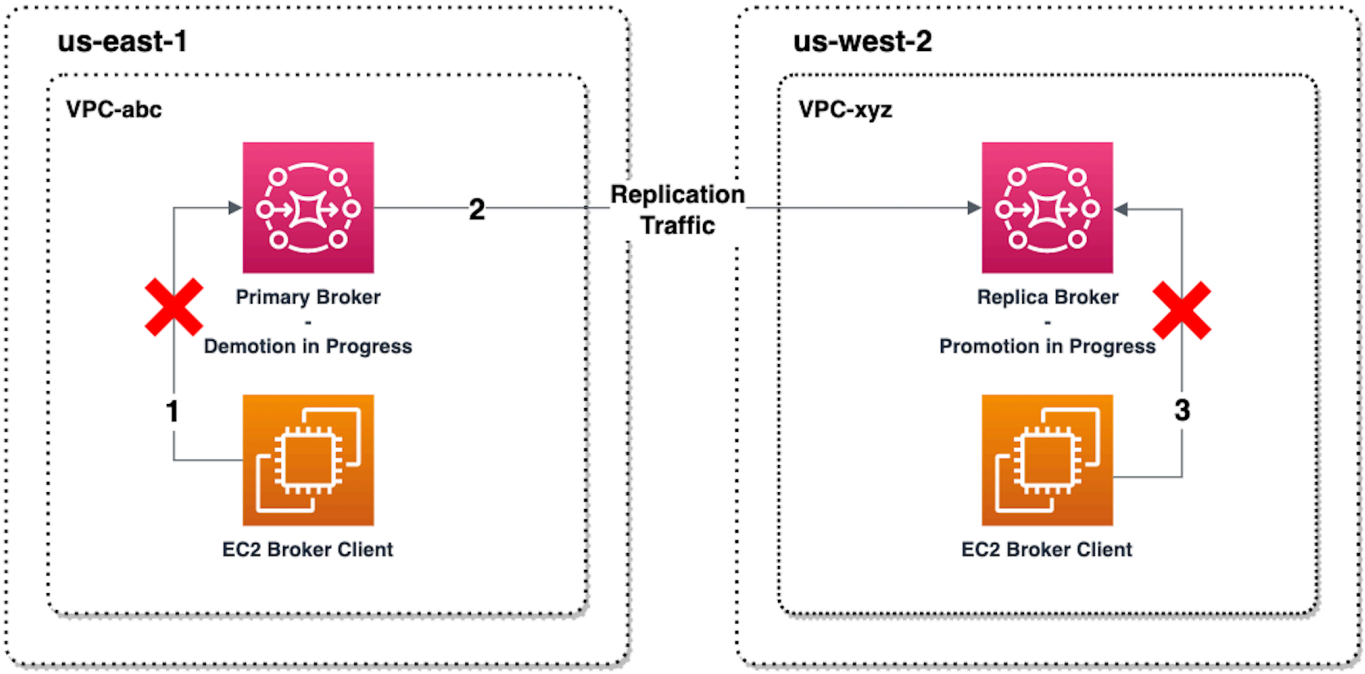

The following diagram illustrates the broker status after the switchover process is complete. The original replica broker has now been promoted to the primary broker role and is accepting client connections. The client can produce and consume data from the broker.

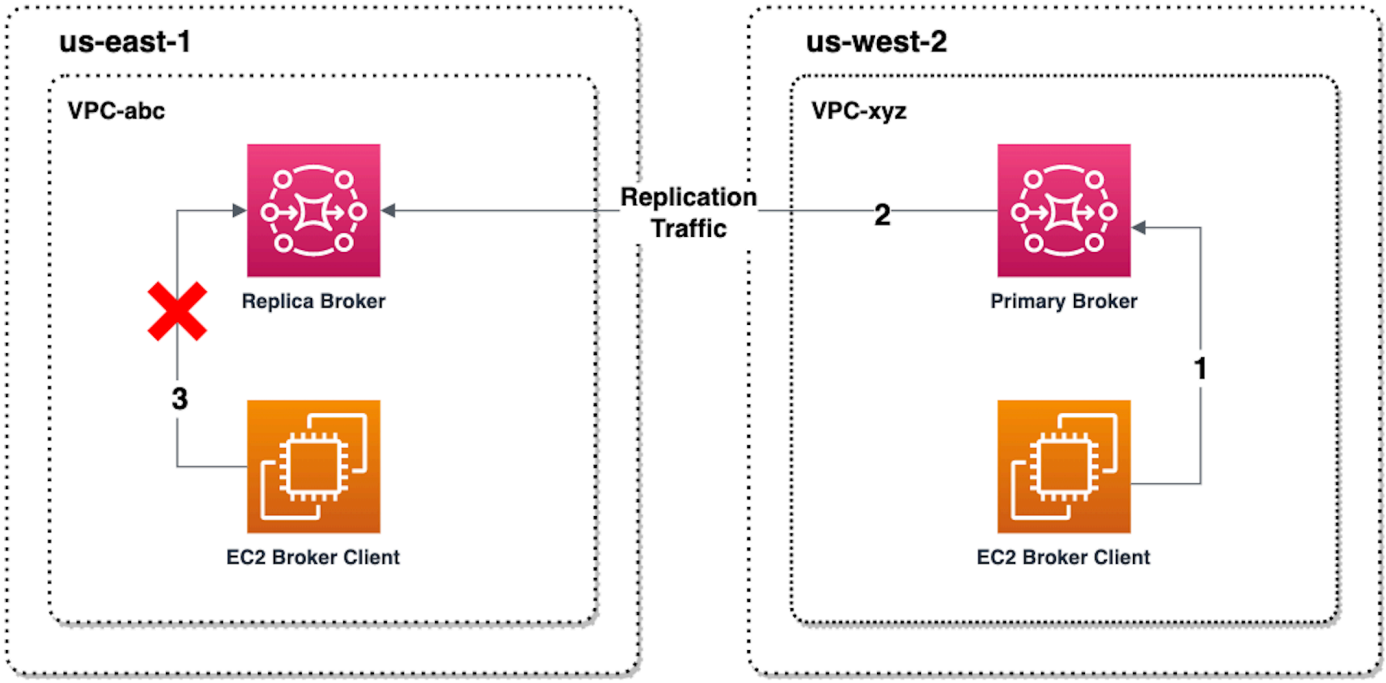

Promote the replica broker using the console

To promote the replica broker using switchover or failover, follow these steps in the Amazon MQ console.

> **Note**
>
> You cannot initiate switchover or failover on a primary broker.

1. Switch to the region for your replica broker. From your Brokers table, select the existing replica broker you will promote to primary.
2. On the **Broker details page**, do the following:
    1. Select **Promote replica**.
    2. In the pop up window, chose *Switchover* or *Failover*.
    3. Type “confirm” in the text box to confirm your choice.
    4. Choose **Confirm**.

After initiating failover, the broker status changes to *Failover in progress*. The blue progress bar at the top of the Brokers page becomes green when failover is complete.

> **Note**
>
> The configuration is only replicated at the time the replicat broker is created. Any update afterwards is not replicated.

Cross-Region data replication metrics in Amazon CloudWatch

The Amazon MQ for ActiveMQ cross-Region data replication feature offers metrics for maintaining the reliability, availability, and performance of your primary and replica brokers. During the replication process, a replica broker in a secondary Region receives asynchronously replicated data from the primary broker in the primary Region. If the primary broker in the primary Region fails, you can promote the replica broker in the secondary Region to primary by initiating a switchover or failover. For instructions on viewing metrics in Amazon CloudWatch, see Accessing CloudWatch metrics for Amazon MQ.

CRDR timestamps

The following timestamps describe how the metrics found in Amazon CloudWatch are calculated. There are five timestamps in the data replication process:

- Time of current observation (TCO): The current instant in time.
- Time of creation (TC): The instant in time an event was created on the replication queue by the primary broker. Available on both primary and replica brokers.
- Time of delivery (TD): The instant in time an event was successfully delivered to the replica broker. Only available on replica brokers.
- Time of processing (TP): The instant in time an event was successfully processed by the replica broker. Only available on replica brokers.
- Time of acknowledgement (TA): The instant in time an event was successfully acknowledged by the primary broker. Only available on primary brokers.

Estimate switchover/failover performance with CRDR CloudWatch metrics

Amazon MQ enables metrics for your broker by default. You can view your broker metrics by accessing the Amazon CloudWatch console, or by using the CloudWatch API. The following metrics are useful for understanding the replication and switchover/failover performance of your CRDR brokers:

Amazon MQ CloudWatch metric	Reason for CRDR use
`TotalReplicationLag`	The estimated time between TA and TC of the last unacknowledged event on the primary broker.
`ReplicationLag`	The estimated time between TP and TC of the last unacknowledged event on the replica broker.
`PrimaryWaitTime`	The estimated time between TCO and TC of the last

Amazon MQ CloudWatch metric	Reason for CRDR use	
	processed event on the primary broker.	
`ReplicaWaitTime`	The estimated time between TCO and TP of the last processed event on the replica broker.	
`QueueSize`	The total number of unacknowledged events in the replication queue on the primary broker.	

`TotalReplicationLag` and `ReplicationLag` describe the delayed replication between the primary and replica brokers. The two metrics can also be used to estimate the time until the ongoing switchover or failover operation complete.

`PrimaryWaitTime` and `ReplicaWaitTime` can be used to identify any ongoing issues with the replication process. If the value of the metric is constantly growing, this can indicate the replication process is degraded or paused. Slow replication may happen due issues like to network partitioning, broker starts, and long recovery.

Quotas in Amazon MQ for ActiveMQ

This topic lists quotas within Amazon MQ. Many of the following quotas can be changed for specific AWS accounts. To request an increase for a limit, see AWS Service Quotas in the *Amazon Web Services General Reference*. Updated limits will not be visible even after the limit increase has been applied. For more information on viewing current connection limits in Amazon CloudWatch, see Monitoring Amazon MQ brokers using Amazon CloudWatch.

> **Note**
>
> For quotas on Amazon MQ for RabbitMQ, see Quotas in Amazon MQ for RabbitMQ.

Topics

Brokers

The following table lists quotas related to Amazon MQ for ActiveMQ brokers.

Limit	Description
Broker name	• Must be unique in the broker region and your AWS account. • Must be 1-50 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric characters, dashes, periods, underscores, and tildes (`-` `.` `_` `~`).
Number of brokers, per region	50
Wire-level connections per protocol for smaller broker	300 for `mq.*.micro` instance type brokers.
Wire-level connections per protocol for larger broker	2,000 for `mq.*.*large` instance type brokers.
Number of network connectors	20
Security groups per broker	5

Limit	Description
ActiveMQ destinations (queues, and topics) monitored in CloudWatch	CloudWatch monitors only the first 1000 destinations.
RabbitMQ destinations (queues) monitored in CloudWatch	CloudWatch monitors only the first 500 destinations, ordered by number of consumers .
Tags per broker	50

Configurations

The following table lists quotas related to Amazon MQ for ActiveMQ configurations.

Limit	Description
Configuration name	• Must be 1-150 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric character s, dashes, periods, underscores, and tildes (- . _ ~).
Revisions per configuration	300

Users

The following table lists quotas related to Amazon MQ for ActiveMQ broker users.

Limit	Description
Username	• Must be 1-100 characters long.

Limit	Description
	• Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~). • Must not contain commas (,).
Password	• Must be 12-250 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Must contain at least 4 unique characters. • Must not contain commas (,).
Users per broker (simple auth)	250
Groups per user (simple auth)	20

Data Storage

The following table lists quotas related to Amazon MQ for ActiveMQ data storage.

Limit	Description
Storage capacity per smaller broker	20 GB for `mq.*.micro` instance type brokers. For more information regarding Amazon MQ instance types, see Broker instance types.
Storage capacity per larger broker	200 GB for `mq.*.*large` instance type brokers. For more information regarding A

Limit	Description
	mazon MQ instance types, see Broker instance types.
Job scheduler usage limit per broker backed by Amazon EBS	50 GB. For more information about job scheduler usage, see `JobSchedulerUsage` in the *Apache ActiveMQ API Documentation*.
Temporary storage capacity per smaller broker.	5 GB for `mq.*.micro` instance type brokers.
Temporary storage capacity per larger broker.	50 GB for `mq.*.*large` instance type brokers.

API Throttling

The following throttling quotas are aggregated per AWS account, *across all Amazon MQ APIs* to maintain service bandwidth. For more information about Amazon MQ APIs, see the *Amazon MQ REST API Reference*.

> **Important**
>
> These quotas don't apply to Amazon MQ for ActiveMQ or Amazon MQ for RabbitMQ broker messaging APIs. For example, Amazon MQ doesn't throttle the sending or receiving of messages.

API burst limit	API rate limit
100	15

Working with Amazon MQ for RabbitMQ

Amazon MQ makes it easy to create a message broker with the computing and storage resources that fit your needs. You can create, manage, and delete brokers using the AWS Management Console, Amazon MQ REST API, or the AWS Command Line Interface.

This section describes the basic elements of a message broker for ActiveMQ and RabbitMQ engine types, lists available Amazon MQ broker instance types and their statuses, and provides an overview of broker architecture and configuration options.

To learn about Amazon MQ REST APIs, see the *Amazon MQ REST API Reference*.

Topics

RabbitMQ engine

This section describes the basic elements of a RabbitMQ broker and its supported plugins, and provides an overview of RabbitMQ broker architecture options on Amazon MQ.

Topics

Basic elements

This section introduces key concepts essential to understanding RabbitMQ on Amazon MQ.

Topics

- Broker defaults
- Broker instance types
- Configurations
- User
- Plugins
- Policies

Broker

A *broker* is a message broker environment running on Amazon MQ. It is the basic building block of Amazon MQ. The combined description of the broker instance *class* (`m5`, `t3`) and *size* (`large`, `micro`) is a *broker instance type* (for example, `mq.m5.large`). For more information, see Broker instance types.

- A *single-instance broker* is comprised of one broker in one Availability Zone behind a Network Load Balancer (NLB) The broker communicates with your application and with an Amazon EBS storage volume.
- A *cluster deployment* is a logical grouping of three RabbitMQ broker nodes behind a Network Load Balancer, each sharing users, queues, and a distributed state across multiple Availability Zones (AZ).

For more information, see Broker architecture.

You can enable *automatic minor version upgrades* to new minor versions of the broker engine, as new versions of the RabbitMQ engine are released. Automatic upgrades occur during the *maintenance window* defined by the day of the week, the time of day (in 24-hour format), and the time zone (UTC by default).

Supported protocols

You can access your RabbitMQ brokers by using any programming language that RabbitMQ supports and by enabling TLS for the following protocols:

- AMQP (0-9-1)

Listener ports

Amazon MQ managed RabbitMQ brokers support the following listener ports for application-level connectivity via `amqps`, as well as client connections using the RabbitMQ web console and the management API.

- Listener port `5671` - Used for connections made via the secure AMQP URL. For example, given a broker with broker ID `b-c8352341-ec91-4a78-ad9c-a43f23d325bb`, deployed in the `us-west-2` region, the following is the broker's full `amqp` URL: `b-c8352341-ec91-4a78-ad9c-a43f23d325bb.mq.us-west-2.amazonaws.com:5671`.
- Listener ports `443` and `15671` - Both listener ports can be used interchangably to access a broker via the RabbitMQ web console or the mangement API.

Attributes

A RabbitMQ broker has several attributes:

- A name. For example, `MyBroker`.
- An ID. For example, `b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`.
- An Amazon Resource Name (ARN). For example, `arn:aws:mq:us-east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`.
- A RabbitMQ web console URL. For example, `https://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com`.

  For more information, see RabbitMQ web console in the RabbitMQ documentation.
- A secure AMQP endpoint. For example, `amqps://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-east-2.amazonaws.com`.

For a full list of broker attributes, see the following in the *Amazon MQ REST API Reference*:

- REST Operation ID: Broker
- REST Operation ID: Brokers
- REST Operation ID: Broker Reboot

Broker defaults

When you create an Amazon MQ for RabbitMQ broker, Amazon MQ applies a default set of broker policies and vhost limits to optimize your broker's performance. Amazon MQ applies vhost limits only to the default (/) vhost. Amazon MQ will not apply default policies to newly created vhosts. We recommend keeping these defaults for all new and existing brokers. However, you can modify, override, or delete these defaults at any time.

Amazon MQ creates policies and limits based on the instance type and broker deployment mode that you choose when you create your broker. The default policies are named according to the deployment mode, as follows:

- **Single-instance** – `AWS-DEFAULT-POLICY-SINGLE-INSTANCE`
- **Cluster deployment** – `AWS-DEFAULT-POLICY-CLUSTER-MULTI-AZ`

For single-instance brokers, Amazon MQ sets the policy priority value to `0`. To override the default priority value, you can create your own custom policies with higher priority values. For cluster deployments, Amazon MQ sets the priority value to 1 for broker defaults. To create your own custom policy for clusters, assign a priority value greater than 1.

> **Note**
>
> In cluster deployments, `ha-mode` and `ha-sync-mode` broker policies are required for classic mirroring and high availability (HA).
> If you delete the default `AWS-DEFAULT-POLICY-CLUSTER-MULTI-AZ` policy, Amazon MQ uses the `ha-all-AWS-OWNED-DO-NOT-DELETE` policy with a priority value of `0`. This ensures that the required `ha-mode` and `ha-sync-mode` policies are still in effect. If you create your own custom policy, Amazon MQ automatically appends `ha-mode` and `ha-sync-mode` to your policy definitions.

Topics

Policy and limit descriptions

The following list describes the default policies and limits that Amazon MQ applies to a newly created broker. The values for `max-length`, `max-queues`, and `max-connections` vary based on your broker's instance type and deployment mode. These values are listed in the Recommended default values section.

- **`queue-mode: lazy`** (policy) – Enables lazy queues. By default, queues keep an in-memory cache of messages, enabling the broker to deliver messages to consumers as fast as possible. This can lead to the broker running out of memory and raising a high-memory alarm. Lazy queues attempt to move messages to disk as early as is practical. This means that fewer messages are kept in memory under normal operating conditions. Using lazy queues, Amazon MQ for RabbitMQ can support much larger messaging loads and longer queues. Note that for certain use cases, brokers with lazy queues might perform marginally slower. This is because messages are moved from disk to broker, as opposed to delivering messages from an in-memory cache.

  > **Deployment modes**
  >
  > Single-instance, cluster

- **`max-length: number-of-messages`** (policy) – Sets a limit for the number of messages in a queue. In cluster deployments, the limit prevents paused queue synchronization in cases such as broker reboots, or following a maintenance window.

  > **Deployment modes**
  >
  > Cluster

- **`overflow: reject-publish`** (policy) – Enforces queues with a `max-length` policy to reject new messages after the number of messages in the queue reaches the `max-length` value. To ensure that messages aren't lost if a queue is in an overflow state, client applications that publish messages to the broker must implement publisher confirms. For information about implementing publisher confirms, see Publisher Confirms on the RabbitMQ website.

  > **Deployment modes**
  >
  > Cluster

- **`max-queues:` *`number-of-queues-per-vhost`*** (vhost limit) – Sets the limit for the number of queues in a broker. Similar to the `max-length` policy definition, limiting the number of queues in cluster deployments prevents paused queue synchronization following broker reboots or maintenance windows. Limiting queues also prevents excessive amounts of CPU usage for maintaining queues.

  > **Deployment modes**
  >
  > Single-instance, cluster

- **`max-connections:` *`number-of-connections-per-vhost`*** (vhost limit) – Sets the limit for the number of client connections to the broker. Limiting the number of connections according to the recommended values prevents excessive broker memory usage, which could result in the broker raising a high memory alarm and pausing operations.

  > **Deployment modes**
  >
  > Single-instance, cluster

Recommended default values

> **Note**
>
> The `max-length` and `max-queue` default limits are tested and evaluated based on an average message size of 5 kB. If your messages are significantly larger than 5 kB, you will need to adjust and reduce the `max-length` and `max-queue` limits.

The following table lists the default limit values for a newly created broker. Amazon MQ applies these values according to the broker's instance type and deployment mode.

Instance type	Deployment mode	`max-length`	`max-queues`	`max-connections`
t3.micro	Single-instance	N/A	500	500
m5.large	Single-instance	N/A	20,000	4,000

Instance type	Deployment mode	`max-length`	`max-queues`	`max-connections`
	Cluster	8,000,000	4,000	15,000
m5.xlarge	Single-instance	N/A	30,000	8,000
	Cluster	9,000,000	5,000	20,000
m5.2xlarge	Single-instance	N/A	60,000	15,000
	Cluster	10,000,000	6,000	40,000
m5.4xlarge	Single-instance	N/A	150,000	30,000
	Cluster	12,000,000	10,000	100,000

Broker instance types

> **Important**
>
> You cannot downgrade a broker from an `mq.m5.` instance type to a `mq.t3.micro` instance type.

> **Important**
>
> You currently cannot create `t2.micro`, `m4.large`, or `m5.*` brokers in **euc1-az2** availability zones.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.t3.micro`	2	1	Low	Use the `mq.t3.micro` instance type for basic evaluation

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
				of Amazon MQ. This instance type (single-instance brokers only) qualifies for the AWS Free Tier. **Important** The `mq.t3.micro` instance type does not support cluster deployment.
`mq.m5.large`	2	8	High	Use the `mq.m5.large` instance for regular development, testing, and production workloads.

Instance Type	vCPU	Memory (GiB)	Network Performance	Notes
`mq.m5.xlarge`	4	16	High	Use the `mq.m5.xlarge` , `mq.m5.2xlarge` , and `mq.m5.4xlarge` instance types for regular development, testing and production workloads that require high throughput.
`mq.m5.2xlarge`	8	32	High	
`mq.m5.4xlarge`	16	64	High	

Configurations

A *configuration* contains all of the settings for your RabbitMQ broker in Cuttlefish format. You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers

> **Important**
>
> Making changes to a configuration does *not* apply the changes to the broker immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle. Currently, you can't delete a configuration.

For information about creating, editing, and managing configurations, see the following:

- Creating and applying broker configurations
- RabbitMQ Broker Configurations

To keep track of the changes you make to your configuration, you can create *configuration revisions*. For more information, see Creating and applying broker configurations.

Attributes

A broker configuration has several attributes, for example:

- A name (`MyConfiguration`)
- An ID (`c-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`)
- An Amazon Resource Name (ARN) (`arn:aws:mq:us-east-2:123456789012:configuration:c-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`)

For a full list of configuration attributes, see the following in the *Amazon MQ REST API Reference*:

- REST Operation ID: Configuration
- REST Operation ID: Configurations

For a full list of configuration revision attributes, see the following:

- REST Operation ID: Configuration Revision
- REST Operation ID: Configuration Revisions

User

Every AMQP 0-9-1 client connection has an associated user which must be authenticated. Each client connection also targets a virtual host (vhost) for which the user must have a set of permissions. A user may have permission to **configure**, **write** to, and **read** from queues and exchanges in a vhost. User credentials, and the target vhost are specified at the time the connection is established.

When you first create an Amazon MQ for RabbitMQ broker, Amazon MQ uses the sign-in credentials you provide to create a RabbitMQ user with the `administrator` tag. You can then add and manage users via the RabbitMQ management API or the RabbitMQ web console. You can also use the RabbitMQ web console or the management API to set or modify user permissions and tags.

> **Note**
>
> RabbitMQ users will not be stored or displayed via the Amazon MQ Users API.

> **Important**
>
> Amazon MQ for RabbitMQ does not support the username "guest", and will delete the default guest account when you create a new broker. Amazon MQ will also periodically delete any customer created account called "guest".

To create a new user with the RabbitMQ management API, use the following API endpoint and request body. Replace *username* and *password* with your new sign-in credentials.

```
PUT /api/users/username HTTP/1.1

{"password":"password","tags":"administrator"}
```

> **Important**
>
> - Do not add personally identifiable information (PII) or other confidential or sensitive information in broker usernames. Broker usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.
> - If you've forgotten the admin password you set while creating the broker, you cannot reset your credentials. If you've created multiple administrators, you can log in using another admin user and reset or recreate your credentials. If you have only one admin user, you must delete the broker and create a new one with new credentials. We recommend consuming or backing up messages before deleting the broker.

The `tags` key is mandatory, and is a comma-separated list of tags for the user. Amazon MQ supports `administrator`, `management`, `monitoring`, and `policymaker` user tags.

You can set permissions for an individual user by using the following API endpoint and request body. Replace *vhost* and *username* with your information. For the default vhost `/`, use `%2F`.

```
POST /api/permissions/vhost/username HTTP/1.1

{"configure":".*","write":".*","read":".*"}
```

> **Note**
>
> The `configure`, `read`, and `write` keys are all mandatory.

By using the wildcard `.*` value, this operation will grant read, write, and configure permissions for all queues in the specified vhost to the user. For more information about managing users via the RabbitMQ management API, see RabbitMQ Management HTTP API.

Plugins

Amazon MQ for RabbitMQ supports the RabbitMQ management plugin which powers the management API and the RabbitMQ web console. You can use the web console and the management API to create and manage broker users and policies.

In addition to the management plugin, Amazon MQ for RabbitMQ also supports the following plugins.

Topics

- Shovel plugin
- Federation plugin
- Consistent Hash exchange plugin

Shovel plugin

Amazon MQ managed brokers support RabbitMQ shovel, allowing you to move messages from queues and exchanges on one broker instance to another. You can use shovel to connect loosely coupled brokers and distribute messages away from nodes with heavier message loads.

Amazon MQ managed RabbitMQ brokers support dynamic shovels. Dynamic shovels are configured using runtime parameters, and can be started and stopped at any time programatically by a client connection. For example, using the RabbitMQ management API, you can create a PUT request to the following API endpoint to configure a dynamic shovel. In the example, `{vhost}` can be

replaced by the name of the broker's vhost, and `{name}` replaced by the name of the new dynamic shovel.

```
/api/parameters/shovel/{vhost}/{name}
```

In the request body, you must specify either a queue or an exchange but not both. This example below configures a dynamic shovel between a local queue specified in `src-queue` and a remote queue defined in `dest-queue`. Similairly, you can use `src-exchange` and `dest-exchange` parameters to configure a shovel between two exchanges.

```
{
  "value": {
    "src-protocol": "amqp091",
    "src-uri":  "amqp://localhost",
    "src-queue":  "source-queue-name",
    "dest-protocol": "amqp091",
    "dest-uri": "amqps://b-c8352341-ec91-4a78-ad9c-a43f23d325bb.mq.us-
west-2.amazonaws.com:5671",
    "dest-queue": "destination-queue-name"
  }
}
```

> **Important**
>
> You cannot configure shovel between queues or exchanges if the shovel destination is a private broker. You can only configure shovel between queues or exchanges in public brokers, or between a source in a private broker, and a destination in a public broker.

For more information about using dynamic shovels, see RabbitMQ dynamic shovel plugin.

> **Note**
>
> Amazon MQ does not support using static shovels.

Federation plugin

Amazon MQ supports federated exchanges and queues. With federation, you can replicate the flow of messages between queues, exchanges and consumers on separate brokers. Federated queues

and exchanges use point-to-point links to connect to peers in other brokers. While federated exchanges, by default, route messages once, federated queues can move messages any number of times as needed by consumers.

You can use federation to allow a *downstream* broker to consume a message from an exchange or a queue on an *upstream*. You can enable federation on downstream brokers by using the RabbitMQ web console or the management API.

> **Important**
>
> You cannot configure federation if the upstream queue or exchange is in a private broker. You can only configure federation between queues or exchanges in public brokers, or between an upstream queue or exchange in a public broker, and a downstream queue or exchange in a private broker.

For example, using the management API, you can configure federation by doing the following.

- Configure one or more upstreams that define federation connections to other nodes. You can define federation connections by using the RabbitMQ web console or the management API. Using the management API, you can create a POST request to `/api/parameters/federation-upstream/%2f/my-upstream` with the following request body.

  ```
  {"value":{"uri":"amqp://server-name","expires":3600000}}
  ```

- Configure a policy to enable your queues or exchanges to become federated. You can configure policies by using the RabbitMQ web console, or the management API. Using the management API, you can create a POST request to `/api/policies/%2f/federate-me` with the following request body.

  ```
  {"pattern":"^amq\.", "definition":{"federation-upstream-set":"all"}, "apply-to":"exchanges"}
  ```

  > **Note**
  >
  > The request body assumes exchanges on the server are named beginning with amq. Using regular expression `^amq\.` will ensure that federation is enabled for all exchanges

whose names begin with "amq." The exchanges on your RabbitMQ server can be named differently.

For more information about configuring the federation plugin, see RabbitMQ federation plugin.

Consistent Hash exchange plugin

By default, Amazon MQ for RabbitMQ supports the Consistent Hash exchange type plugin. Consistent Hash exchanges route messages to queues based on a hash value calculated from the *routing key* of a message. Given a reasonably even routing key, Consistent Hash exchanges can distribute messages between queues reasonably evenly.

For queues bound to a Consistent Hash exchange, the binding key is a number-as-a-string that determines the *binding weight* of each queue. Queues with a higher binding weight will receive a proportionally higher distribution of messages from the Consistent Hash exchange to which they are bound. In a Consistent Hash exchange topology, publishers can simply publish messages to the exchange, but consumers must be explicitly configured to consume messages from specific queues.

For more information about Consistent Hash exchanges, see RabbitMQ Consistent Hash Exchange Type on the GitHub website.

Policies

You can apply custom policies and limits with Amazon MQ recommended default values. If you have deleted the recommended default policies and limits, and want to re-create them, or you have created additional vhosts and want to apply the default policies and limits to your new vhosts, you can use the following steps.

Important

To perform the following steps, you must have an Amazon MQ for RabbitMQ broker user with administrator permissions. You can use the administrator user created when you first created the broker, or another user that you might have created afterwards. The following table provides the necessary administrator user tag and permissions as regular expression (regexp) patterns.

Tags	Read regexp	Configure regexp	Write regexp
`administrator`	`.*`	`.*`	`.*`

For more information about creating RabbitMQ users and managing user tags and permissions, see User.

To apply default policies and virtual host limits using the RabbitMQ web console

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**.
3. From the list of brokers, choose the name of the broker to which you want to apply the new policy.
4. On the broker details page, in the **Connections** section, choose the **RabbitMQ web console** URL. The RabbitMQ web console opens in a new browser tab or window.
5. Log in to the RabbitMQ web console with your broker administrator user name and password.
6. In the RabbitMQ web console, at the top of the page, choose **Admin**.
7. On the **Admin** page, in the right navigation pane, choose **Policies**.
8. On the **Policies** page, you can see a list of the broker's current **User policies**. Below **User policies**, expand **Add / update a policy**.
9. To create a new broker policy, under **Add / update a policy**, do the following:

   a. For **Virtual host**, choose the name of the vhost to which you want to attach the policies from the dropdown list. To choose the default vhost, choose **/**.

   > **Note**
   >
   > If you have not created additional vhosts, the **Virtual host** option is not shown in the RabbitMQ console, and the policies are applied only to the default vhost.

   b. For **Name**, enter a name for your policy, for example, **`policy-defaults`**.

   c. For **Pattern**, enter the regexp pattern **`.*`** so that the policy matches all queues on the broker.

d. For **Apply to**, choose **Exchanges and queues** from the dropdown list.

e. For **Priority**, enter an integer greater than all other policies applied to the vhost. You can apply exactly one set of policy definitions to RabbitMQ queues and exchanges at any given time. RabbitMQ chooses the matching policy with the highest priority value. For more information about policy priorities and how to combine policies, see Policies in the RabbitMQ Server Documentation.

f. For **Definition**, add the following key-value pairs:

- **`queue-mode`**=**`lazy`**. Choose **String** from the dropdown list.
- **`overflow`**=**`reject-publish`**. Choose **String** from the dropdown list.

> **Note**
>
> Does not apply to single-instance brokers.

- **`max-length`**=***`number-of-messages`***. Replace *`number-of-messages`* with the Amazon MQ recommended value according to the broker's instance size and deployment mode, for example, **`8000000`** for an `mq.m5.large` cluster. Choose **Number** from the dropdown list.

> **Note**
>
> Does not apply to single-instance brokers.

g. Choose **Add / update policy**.

10. Confirm that the new policy appears in the list of **User policies**.

> **Note**
>
> For cluster brokers, Amazon MQ automatically applies the `ha-mode: all` and `ha-sync-mode: automatic` policy definitions.

11. From the right navigation pane, choose **Limits**.
12. On the **Limits** page, you can see a list of the broker's current **Virtual host limits**. Below **Virtual host limits**, expand **Set / update a virtual host limit**.
13. To create a new vhost limit, under **Set / update a virtual host limit**, do the following:

a. For **Virtual host**, choose the name of the vhost to which you want to attach the policies from the dropdown list. To choose the default vhost, choose **/**.

b. For **Limit**, choose **max-connections** from the dropdown options.

c. For **Value**, enter the Amazon MQ recommended value according to the broker's instance size and deployment mode, for example, **`15000`** for an `mq.m5.large` cluster.

d. Choose **Set / update limit**.

e. Repeat the steps above, and for **Limit**, choose **max-queues** from the dropdown options.

14. Confirm that the new limits appear in the list of **Virtual host limits**.

To apply default policies and virtual host limits using the RabbitMQ management API

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**.
3. From the list of brokers, choose the name of the broker to which you want to apply the new policy.
4. On the broker's page, in the **Connections** section, note the **RabbitMQ web console** URL. This is the broker endpoint that you use in an HTTP request.
5. Open a new terminal or command line window of your choice.
6. To create a new broker policy, enter the following `curl` command. This command assumes a queue on the default `/` vhost, which is encoded as `%2F`. To apply the policy to another vhost, replace `%2F` with the vhost's name.

> **Note**
>
> Replace *`username`* and *`password`* with your administrator sign-in credentials. Replace *`number-of-messages`* with the Amazon MQ recommended value according to the broker's instance size and deployment mode. Replace *`policy-name`* with a name for your policy. Replace *`broker-endpoint`* with the URL that you noted previously.

```
curl -i -u username:password -H "content-type:application/json" -XPUT \
-d '{"pattern":".*", "priority":1, "definition":{"queue-mode":lazy,
 "overflow":"reject-publish", "max-length":"number-of-messages"}}' \
broker-endpoint/api/policies/%2F/policy-name
```

7. To confirm that the new policy is added to your broker's user policies, enter the following `curl` command to list all broker policies.

```
curl -i -u username:password broker-endpoint/api/policies
```

8. To create a new `max-connections` virtual host limits, enter the following `curl` command. This command assumes a queue on the default `/` vhost, which is encoded as `%2F`. To apply the policy to another vhost, replace `%2F` with the vhost's name.

> **Note**
>
> Replace *`username`* and *`password`* with your administrator sign-in credentials. Replace *`max-connections`* with the Amazon MQ recommended value according to the broker's instance size and deployment mode. Replace the broker endpoint with the URL that you noted previously.

```
curl -i -u username:password -H "content-type:application/json" -XPUT \
-d '{"value":"number-of-connections"}' \
broker-endpoint/api/vhost-limits/%2F/max-connections
```

9. To create a new `max-queues` virtual host limit, repeat the previous step, but modify the curl command as shown in the following.

```
curl -i -u username:password -H "content-type:application/json" -XPUT \
-d '{"value":"number-of-queues"}' \
broker-endpoint/api/vhost-limits/%2F/max-queues
```

10. To confirm that the new limits are added to your broker's virtual host limits, enter the following `curl` command to list all broker virtual host limits.

```
curl -i -u username:password broker-endpoint/api/vhost-limits
```

Broker architecture

RabbitMQ brokers can be created as *single-instance brokers* or in a *cluster deployment*. For both deployment modes, Amazon MQ provides high durability by storing its data redundantly.

You can access your RabbitMQ brokers by using any programming language that RabbitMQ supports and by enabling TLS for the following protocols:

- AMQP (0-9-1)

Topics

- Single-instance broker
- Cluster deployment for high availability

Single-instance broker

A *single-instance broker* is comprised of one broker in one Availability Zone behind a Network Load Balancer (NLB). The broker communicates with your application and with an Amazon EBS storage volume. Amazon EBS provides block level storage optimized for low-latency and high throughput.

Using an Network Load Balancer ensures that your Amazon MQ for RabbitMQ broker endpoint remains unchanged if the broker instance is replaced during a maintenance window or because of underlying Amazon EC2 hardware failures. An Network Load Balancer allows your applications and users to continue to use the same endpoint to connect to the broker.

The following diagram illustrates an Amazon MQ for RabbitMQ single-instance broker.

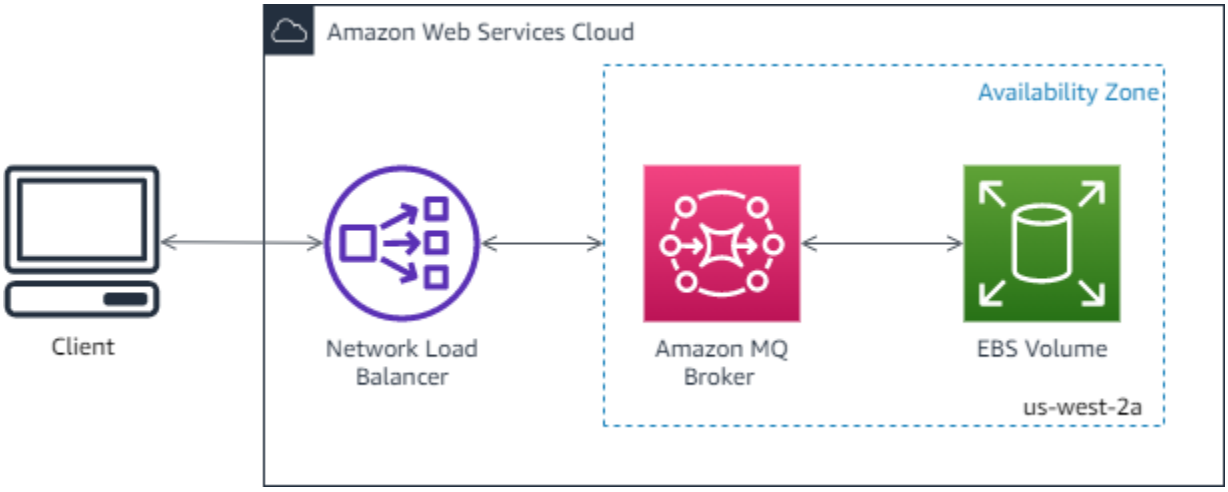

Cluster deployment for high availability

A *cluster deployment* is a logical grouping of three RabbitMQ broker nodes behind a Network Load Balancer, each sharing users, queues, and a distributed state across multiple Availability Zones (AZ).

In a cluster deployment, Amazon MQ automatically manages broker policies to enable classic mirroring across all nodes, ensuring high availability (HA). Each mirrored queue consists of one *main* node and one or more *mirrors*. Each queue has its own main node. All operations for a given queue are first applied on the queue's main node and then propagated to mirrors. Amazon MQ creates a default system policy that sets the `ha-mode` to `all` and `ha-sync-mode` to `automatic`. This ensures that data is replicated to all nodes in the cluster across different Availability Zones for better durability.

> **Note**
>
> During a *maintenance window*, all maintenance to a cluster is performed one node at a time, keeping at least two running nodes at all times. Each time a node is brought down, client connections to that node are severed and need to be re-established. You must ensure that your client code is designed to automatically reconnect to your cluster. For more information about connection recovery, see the section called "Automatically recover from network failures".
> Because Amazon MQ sets `ha-sync-mode: automatic`, during a maintenance window, queues will synchronize when each node re-joins the cluster. Queue synchronization blocks all other queue operations. You can mitigate the impact of queue synchronization during maintenance windows by keeping queues short.

The default policy should not be deleted. If you do delete this policy, Amazon MQ will be automatically recreate it. Amazon MQ will also ensure that HA properties are applied to all other policies that you create on a clustered broker. If you add a policy without the HA properties, Amazon MQ will add them for you. If you add a policy with different high availability properties, Amazon MQ will replace them. For more information about classic mirroring, see Classic mirrored queues.

> **Important**
>
> Amazon MQ does not support quorum queues. Enabling the quorum queue feature flag and creating quorum queues will result in data loss.

The following diagram illustrates a RabbitMQ cluster broker deployment with three nodes in three Availability Zones (AZ), each with its own Amazon EBS volume and a shared state. Amazon EBS provides block level storage optimized for low-latency and high throughput.

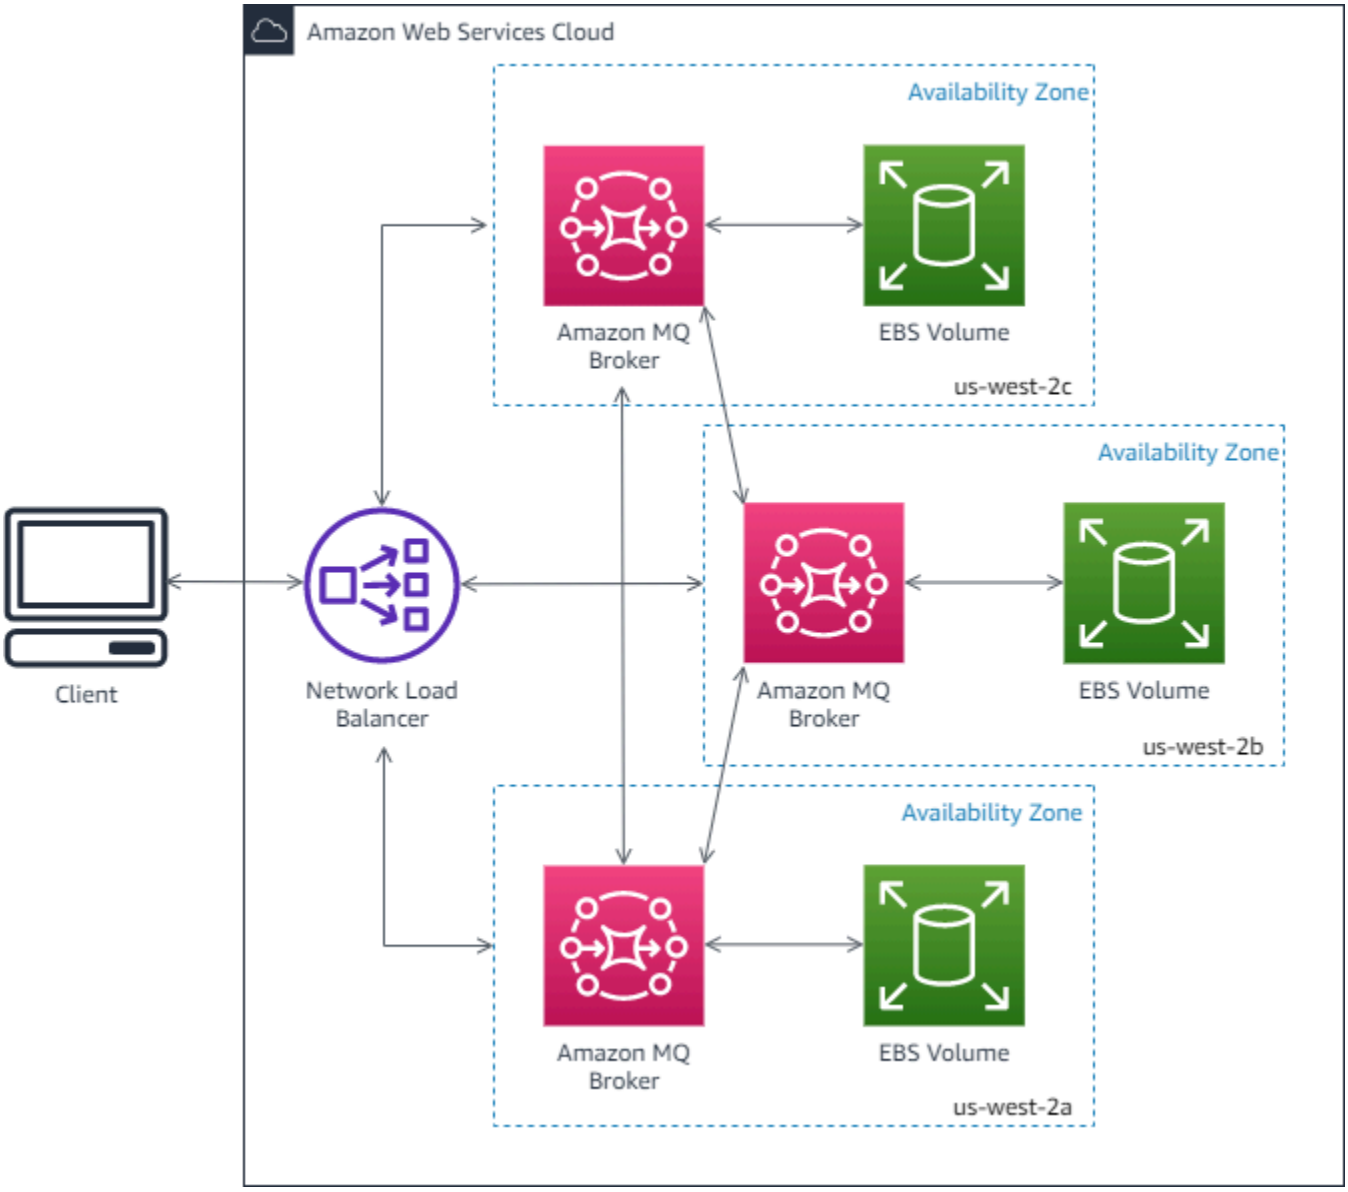

Amazon MQ for RabbitMQ broker configurations

A configuration contains all of the settings for your RabbitMQ broker in Cuttlefish format. You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers.

Topics

- Creating, editing, and applying RabbitMQ broker configurations

- RabbitMQ configuration policies

Creating, editing, and applying RabbitMQ broker configurations

A *configuration* contains all of the settings for your RabbitMQ broker in Cuttlefish format. You can create a configuration before creating any brokers. You can then apply the configuration to one or more brokers

For more information, see the following:

- Configurations
- Amazon MQ broker configuration lifecycle

The following examples show how you can create and apply a RabbitMQ broker configuration using the AWS Management Console.

Topics

- Create a New Configuration
- Create a New Configuration Revision
- Apply a Configuration Revision to Your Broker
- Edit a Configuration Revision

Create a New Configuration

1. Sign in to the Amazon MQ console.
2. On the left, expand the navigation panel and choose **Configurations**.

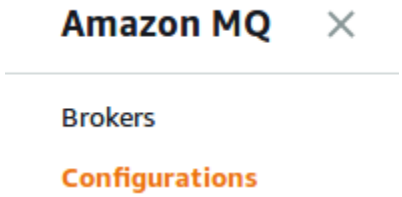

3. On the **Configurations** page, choose **Create configuration**.
4. On the **Create configuration** page, in the **Details** section, type the **Configuration name** (for example, `MyConfiguration`) and select a **Broker engine** version.

To learn more about RabbitMQ engine versions supported by Amazon MQ for RabbitMQ, see the section called "Version management".

5. Choose **Create configuration**.

Create a New Configuration Revision

1. From the configuration list, choose ***MyConfiguration***.

> **Note**
>
> The first configuration revision is always created for you when Amazon MQ creates the configuration.

On the ***MyConfiguration*** page, the broker engine type and version that your new configuration revision uses (for example, **RabbitMQ 3.xx.xx**) are displayed.

2. On the **Configuration details** tab, the configuration revision number, description, and broker configuration in Cuttlefish format are displayed.

> **Note**
>
> Editing the current configuration creates a new configuration revision.

3. Choose **Edit configuration** and make changes to the Cuttlefish configuration.
4. Choose **Save**.

The **Save revision** dialog box is displayed.

5. (Optional) Type `A description of the changes in this revision`.
6. Choose **Save**.

The new revision of the configuration is saved.

> **Important**
>
> Making changes to a configuration does *not* apply the changes to the broker immediately. To apply your changes, you must wait for the next maintenance window

or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.
Currently, you can't delete a configuration.

Apply a Configuration Revision to Your Broker

1. On the left, expand the navigation panel and choose **Brokers**.

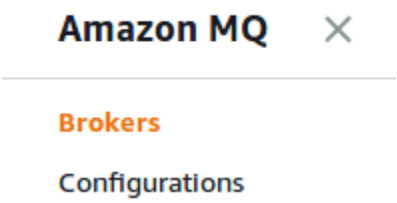

2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Edit**.
3. On the **Edit *MyBroker*** page, in the **Configuration** section, select a **Configuration** and a **Revision** and then choose **Schedule Modifications**.
4. In the **Schedule broker modifications** section, choose whether to apply modifications **During the next scheduled maintenance window** or **Immediately**.

> **Important**
>
> Your broker will be offline while it is being rebooted.

5. Choose **Apply**.

   Your configuration revision is applied to your broker at the specified time.

Edit a Configuration Revision

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Edit**.
3. On the ***MyBroker*** page, choose **Edit**.
4. On the **Edit *MyBroker*** page, in the **Configuration** section, select a **Configuration** and a **Revision** and then choose **Edit**.

> **Note**
>
> Unless you select a configuration when you create a broker, the first configuration revision is always created for you when Amazon MQ creates the broker.

On the ***MyBroker*** page, the broker engine type and version that the configuration uses (for example, **RabbitMQ 3.xx.xx**) are displayed.

5. On the **Configuration details** tab, the configuration revision number, description, and broker configuration in Cuttlefish format are displayed.

> **Note**
>
> Editing the current configuration creates a new configuration revision.

6. Choose **Edit configuration** and make changes to the Cuttlefish configuration.
7. Choose **Save**.

   The **Save revision** dialog box is displayed.
8. (Optional) Type `A description of the changes in this revision`.
9. Choose **Save**.

   The new revision of the configuration is saved.

> **Important**
>
> Making changes to a configuration does *not* apply the changes to the broker immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.
> Currently, you can't delete a configuration.

RabbitMQ configuration policies

Amazon MQ for RabbitMQ now supports creating and applying configurations to your RabbitMQ broker. The default operator policy on each virtual host has the following recommended HA properties:

```
name: default_operator_policy_AWS_managed
pattern: .*
apply-to: all
priority: 0
definition: {
ha-mode: all
ha-sync-mode: automatic
}
```

Changes to the operator policies via the AWS Management Console or Management API are not available by default. You can enable changes by adding the following line to the broker configuration:

```
management.restrictions.operator_policy_changes.disabled=false
```

If you make this change, you are highly encouraged to include the HA properties in your own operator policies. For more information on adding configurations to your broker, see Creating and applying broker configurations.

Managing Amazon MQ for RabbitMQ engine versions

RabbitMQ organizes version numbers according to semantic versioning specification as `X.Y.Z`. In Amazon MQ for RabbitMQ implementations, `X.Y` denotes the major version, and `Z` represents the minor version number. Amazon MQ considers a version change to be major if the major version numbers change. For example, upgrading from version 3.**8** to 3.**9** is considered a *major version upgrade*. A version change is considered minor if only the minor version number changes. For example, upgrading from version 3.8.**23** to 3.8.**26** is considered a *minor version upgrade*.

When you create a new Amazon MQ for RabbitMQ broker, you can specify any supported RabbitMQ engine version. If you use the AWS Management Console to create a broker, Amazon MQ automatically defaults to the latest engine version number. If you use the AWS CLI or the Amazon

MQ API to create a broker, the engine version number is required. If you don't provide a version number, the operation will result in an exception. To learn more, see `create-broker` in the *AWS CLI Command Reference* and `CreateBroker` in the *Amazon MQ REST API Reference*.

> **Important**
>
> Amazon MQ does not support quorum queues or streams. Enabling these feature flag(s) and creating a quorum queue or stream will result in data loss.

> **Important**
>
> - Amazon MQ does not support using structured logging in JSON, introduced in RabbitMQ 3.9
> - Amazon MQ for RabbitMQ recommends new brokers to use the latest supported minor version.
> - RabbitMQ only allows incremental version updates (ex: 3.9.x to 3.10.x). You cannot skip major versions when updating (ex: 3.8.x to 3.10.x).

Amazon MQ for RabbitMQ currently supports the following engine versions:

Major versions	Minor versions
RabbitMQ 3.11	• 3.11.28 (recommended)
RabbitMQ 3.10	• 3.10.25
RabbitMQ 3.9	• 3.9.27
RabbitMQ 3.8	• 3.8.34

The following minor versions are still available for existing brokers, but they are not recommended for new brokers.

Existing minor versions

Major versions	Minor versions
RabbitMQ 3.11	• 3.11.20 • 3.11.16
RabbitMQ 3.10	• 3.10.20 • 3.10.10
RabbitMQ 3.9	• 3.9.24 • 3.9.16 • 3.9.13
RabbitMQ 3.8	• 3.8.30 • 3.8.27 • 3.8.26 • 3.8.23 • 3.8.22 • 3.8.11 • 3.8.6

Major and minor version upgrades

With Amazon MQ, you control when to upgrade your brokers to new versions. When automatic minor version upgrade is activated, Amazon MQ will automatically upgrade your broker engine to new minor versions as they are released and supported by Amazon MQ.

To perform a major version upgrade, you must manually upgrade your broker's engine version number. Minor and major version upgrades occure at the same time as other broker patching operations, during your scheduled maintenance window. If you opt out of automatic minor version upgrades, you can manually upgrade your broker to a new supported minor version by following the same procedure as a major upgrade.

For more information about updating your broker preferences to activate or deactivate minor version upgrades, and manually upgrading your broker, see the section called "Upgrading the engine version".

Single instance brokers will be offline while being rebooted. For cluster brokers, the mirrored queues must be synced during reboot. With longer queues, the queue-sync process can take longer. During the queue-sync process, the queue is unavailable to consumers and producer. When the queue-sync process is complete, the broker becomes available again. To minimize the impact, we recommend upgrading during a low traffic time. For more information on best practices for version upgrades, see Amazon MQ for RabbitMQ best practices.

Listing supported engine versions

You can list all supported minor and major engine versions by using the `describe-broker-instance-options` AWS CLI command.

```
aws mq describe-broker-instance-options
```

To filter the results by engine and instance type use the `--engine-type` and `--host-instance-type` options as shown in the following.

```
aws mq describe-broker-instance-options --engine-type engine-type --host-instance-
type instance-type
```

For example, to filter the results for RabbitMQ, and `mq.m5.large` instance type, replace *`engine-type`* with `RABBITMQ` and *`instance-type`* with `mq.m5.large`.

RabbitMQ tutorials

The following tutorials show how you can configure and use RabbitMQ on Amazon MQ. To learn more about working with supported client libraries in a variety of programming languages such as Node.js, Python, .NET, and more, see RabbitMQ Tutorials in the *RabbitMQ Getting Started Guide*.

Topics

Editing broker preferences

You can edit your broker preferences, such as enabling or disabling CloudWatch logs using the AWS Management Console.

Edit RabbitMQ broker options

1. Sign in to the Amazon MQ console.
2. From the broker list, select your broker (for example, **MyBroker**) and then choose **Edit**.
3. On the **Edit *MyBroker*** page, in the **Specifications** section, select a **Broker engine version** or a **Broker Instance type**.
4. In the **CloudWatch Logs** section, click the toggle button to enable or disable general logs. No other steps are required.

   > **Note**
   >
   > - For RabbitMQ brokers, Amazon MQ automatically uses a Service-Linked Role (SLR) to publish general logs to CloudWatch. For more information, see the section called "Using service-linked roles"
   > - Amazon MQ does not support audit logging for RabbitMQ brokers.

5. In the **Maintenance** section, configure your broker's maintenance schedule:

   To upgrade the broker to new versions as AWS releases them, choose **Enable automatic minor version upgrades**. Automatic upgrades occur during the *maintenance window* defined by the day of the week, the time of day (in 24-hour format), and the time zone (UTC by default).
6. Choose **Schedule modifications**.

   > **Note**
   >
   > If you choose only **Enable automatic minor version upgrades**, the button changes to **Save** because no broker reboot is necessary.

   Your preferences are applied to your broker at the specified time.

Using Python Pika with Amazon MQ for RabbitMQ

The following tutorial shows how you can set up a Python Pika client with TLS configured to connect to an Amazon MQ for RabbitMQ broker. Pika is a Python implementation of the AMQP 0-9-1 protocol for RabbitMQ. This tutorial guides you through installing Pika, declaring a queue, setting up a publisher to send messages to the broker's default exchange, and setting up a consumer to recieve messages from the queue.

Topics

Prerequisites

To complete the steps in this tutorial, you need the following prerequisites:

- An Amazon MQ for RabbitMQ broker. For more information, see Creating an Amazon MQ for RabbitMQ broker.
- Python 3 installed for your operating system.
- Pika installed using Python `pip`. To install Pika, open a new terminal window and run the following.

```
$ python3 -m pip install pika
```

Permissions

For this tutorial, you need at least one Amazon MQ for RabbitMQ broker user with permission to write to, and read from, a vhost. The following table describes the neccessary minimum permissions as regular expression (regexp) patterns.

Tags	Configure regexp	Write regexp	Read regexp
none		`.*`	`.*`

The user permissions listed provide only read and write permissions to the user, without granting access to the management plugin to perform administrative operations on the broker. You can further restrict permissions by providing regexp patterns that limit the user's access to specified queues. For example, if you change the read regexp pattern to `^[hello world].*`, the user will only have permission to read from queues that start with `hello world`.

For more information about creating RabbitMQ users and managing user tags and permissions, see User.

Step one: Create a basic Python Pika client

To create a Python Pika client base class that defines a constructor and provides the SSL context necessary for TLS configuration when interacting with an Amazon MQ for RabbitMQ broker, do the following.

1. Open a new terminal window, create a new directory for your project, and navigate to the directory.

```
$ mkdir pika-tutorial
$ cd pika-tutorial
```

2. Create a new file, `basicClient.py`, that contains the following Python code.

```
import ssl
import pika

class BasicPikaClient:

    def __init__(self, rabbitmq_broker_id, rabbitmq_user, rabbitmq_password,
 region):

        # SSL Context for TLS configuration of Amazon MQ for RabbitMQ
        ssl_context = ssl.SSLContext(ssl.PROTOCOL_TLSv1_2)
        ssl_context.set_ciphers('ECDHE+AESGCM:!ECDSA')
```

```
        url = f"amqps://{rabbitmq_user}:
{rabbitmq_password}@{rabbitmq_broker_id}.mq.{region}.amazonaws.com:5671"
        parameters = pika.URLParameters(url)
        parameters.ssl_options = pika.SSLOptions(context=ssl_context)

        self.connection = pika.BlockingConnection(parameters)
        self.channel = self.connection.channel()
```

You can now define additional classes for your publisher and consumer that inherit from `BasicPikaClient`.

Step two: Create a publisher and send a message

To create a publisher that declares a queue, and sends a single message, do the following.

1. Copy the contents of the following code sample, and save locally as `publisher.py` in the same directory you created in the previous step.

```
from basicClient import BasicPikaClient

class BasicMessageSender(BasicPikaClient):

    def declare_queue(self, queue_name):
        print(f"Trying to declare queue({queue_name})...")
        self.channel.queue_declare(queue=queue_name)

    def send_message(self, exchange, routing_key, body):
        channel = self.connection.channel()
        channel.basic_publish(exchange=exchange,
                              routing_key=routing_key,
                              body=body)
        print(f"Sent message. Exchange: {exchange}, Routing Key: {routing_key},
 Body: {body}")

    def close(self):
        self.channel.close()
        self.connection.close()

if __name__ == "__main__":

    # Initialize Basic Message Sender which creates a connection
    # and channel for sending messages.
```

```
    basic_message_sender = BasicMessageSender(
        "<broker-id>",
        "<username>",
        "<password>",
        "<region>"
    )

    # Declare a queue
    basic_message_sender.declare_queue("hello world queue")

    # Send a message to the queue.
    basic_message_sender.send_message(exchange="", routing_key="hello world queue",
 body=b'Hello World!')

    # Close connections.
    basic_message_sender.close()
```

The `BasicMessageSender` class inherits from `BasicPikaClient` and implements additional methods for delaring a queue, sending a message to the queue, and closing connections. The code sample routes a message to the default exchange, with a routing key equal to the name of the queue.

2. Under `if __name__ == "__main__":`, replace the parameters passed to the `BasicMessageSender` constructor statement with the following information.

   - **`<broker-id>`** – The unique ID that Amazon MQ generates for the broker. You can parse the ID from your broker ARN. For example, given the following ARN, `arn:aws:mq:us-east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`, the broker ID would be `b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9`.
   - **`<username>`** – The username for a broker user with sufficient permissions to write messages to the broker.
   - **`<password>`** – The password for a broker user with sufficient permissions to write messages to the broker.
   - **`<region>`** – The AWS region in which you created your Amazon MQ for RabbitMQ broker. For example, `us-west-2`.

3. Run the following command in the same directory you created `publisher.py`.

```
$ python3 publisher.py
```

If the code runs successfully, you will see the following output in your terminal window.

```
Trying to declare queue(hello world queue)...
Sent message. Exchange: , Routing Key: hello world queue, Body: b'Hello World!'
```

Step three: Create a consumer and recieve a message

To create a consumer that recieves a single message from the queue, do the following.

1. Copy the contents of the following code sample, and save locally as `consumer.py` in the same directory.

```
from basicClient import BasicPikaClient

class BasicMessageReceiver(BasicPikaClient):

    def get_message(self, queue):
        method_frame, header_frame, body = self.channel.basic_get(queue)
        if method_frame:
            print(method_frame, header_frame, body)
            self.channel.basic_ack(method_frame.delivery_tag)
            return method_frame, header_frame, body
        else:
            print('No message returned')

    def close(self):
        self.channel.close()
        self.connection.close()


if __name__ == "__main__":

    # Create Basic Message Receiver which creates a connection
    # and channel for consuming messages.
    basic_message_receiver = BasicMessageReceiver(
        "<broker-id>",
        "<username>",
        "<password>",
        "<region>"
    )
```

```
    # Consume the message that was sent.
    basic_message_receiver.get_message("hello world queue")

    # Close connections.
    basic_message_receiver.close()
```

Similar to the the publisher you created in the previous step, `BasicMessageReciever` inherits from `BasicPikaClient` and implements additional methods for recieving a single message, and closing connections.

2. Under the `if __name__ == "__main__":` statement, replace the parameters passed to the `BasicMessageReciever` constructor with your information.
3. Run the following command in your project directory.

```
$ python3 consumer.py
```

If the code runs successfully, you will see the message body, and headers including the routing key, displayed in your terminal window.

```
<Basic.GetOk(['delivery_tag=1', 'exchange=', 'message_count=0',
 'redelivered=False', 'routing_key=hello world queue'])> <BasicProperties> b'Hello
 World!'
```

Step four: (Optional) Set up an event loop and consume messages

To consume multiple messages from a queue, use Pika's `basic_consume` method and a callback function as shown in the following

1. In `consumer.py`, add the following method definition to the `BasicMessageReceiver` class.

```
def consume_messages(self, queue):
    def callback(ch, method, properties, body):
        print(" [x] Received %r" % body)

    self.channel.basic_consume(queue=queue, on_message_callback=callback,
 auto_ack=True)

    print(' [*] Waiting for messages. To exit press CTRL+C')
    self.channel.start_consuming()
```

2. In `consumer.py`, under `if __name__ == "__main__":`, invoke the `consume_messages` method you defined in the previous step.

```
if __name__ == "__main__":

    # Create Basic Message Receiver which creates a connection and channel for
 consuming messages.
    basic_message_receiver = BasicMessageReceiver(
        "<broker-id>",
        "<username>",
        "<password>",
        "<region>"
    )

    # Consume the message that was sent.
    # basic_message_receiver.get_message("hello world queue")

    # Consume multiple messages in an event loop.
    basic_message_receiver.consume_messages("hello world queue")

    # Close connections.
    basic_message_receiver.close()
```

3. Run `consumer.py` again, and if successful, the queued messages will be displayed in your terminal window.

```
[*] Waiting for messages. To exit press CTRL+C
[x] Received b'Hello World!'
[x] Received b'Hello World!'
...
```

What's next?

- For more information about other supported RabbitMQ client libraries, see RabbitMQ Client Documentation on the RabbitMQ website.

Resolving RabbitMQ paused queue synchronization

In an Amazon MQ for RabbitMQ cluster deployment, messages published to each queue are replicated across three broker nodes. This replication, referred to as *mirroring*, provides high availability (HA) for RabbitMQ brokers. Queues in a cluster deployment consist of a *main* replica on one node and one or more *mirrors*. Every operation applied to a mirrored queue, including enqueuing messages, is first applied to the main queue and then replicated across its mirrors.

For example, consider a mirrored queue replicated across three nodes: the main node (`main`) and two mirrors (`mirror-1` and `mirror-2`). If all messages in this mirrored queue are successfully propagated to all mirrors, then the queue is synchronized. If a node (`mirror-1`) becomes unavailable for an interval of time, the queue is still operational and can continue to enqueue messages. However, for the queue to synchronize, messages published to `main` while `mirror-1` is unavailable must be replicated to `mirror-1`.

For more information about mirroring, see Classic Mirrored Queues on the RabbitMQ website.

Maintenance and queue synchronization

During maintenance windows, Amazon MQ performs all maintenance work one node at a time to ensure that the broker remains operational. As a result, queues might need to synchronize as each node resumes operation. During synchronization, messages that need to be replicated to mirrors are loaded into memory from the corresponding Amazon Elastic Block Store (Amazon EBS) volume to be processed in batches. Processing messages in batches lets queues synchronize faster.

If queues are kept short and messages are small, the queues successfully synchronize and resume operation as expected. However, if the amount of data in a batch approaches the node's memory limit, the node raises a high memory alarm, pausing the queue sync. You can confirm memory usage by comparing the `RabbitMemUsed` and `RabbitMqMemLimit` broker node metrics in CloudWatch. Synchronization can't complete until messages are consumed or deleted, or the number of messages in the batch is reduced.

> **Note**
>
> Reducing the queue synchronization batch size can result in a higher number of replication transactions.

To resolve a paused queue synchronization, follow the steps in this tutorial, which demonstrates applying an `ha-sync-batch-size` policy and restarting the queue sync.

Topics

Prerequisites

For this tutorial, you must have an Amazon MQ for RabbitMQ broker user with administrator permissions. You can use the administrator user created when you first created the broker, or another user that you might have created afterwards. The following table provides the necessary administrator user tag and permissions as regular expression (regexp) patterns.

Tags	Read regexp	Configure regexp	Write regexp
`administrator`	`.*`	`.*`	`.*`

For more information about creating RabbitMQ users and managing user tags and permissions, see User.

Step 1: Apply an `ha-sync-batch-size` policy

The following procedures demonstrate adding a policy that applies to all queues created on the broker. You can use the RabbitMQ web console or the RabbitMQ management API. For more information, see Management Plugin on the RabbitMQ website.

To apply an `ha-sync-batch-size` policy using the RabbitMQ web console

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**.
3. From the list of brokers, choose the name of the broker to which you want to apply the new policy.

4. On the broker's page, in the **Connections** section, choose the **RabbitMQ web console** URL. The RabbitMQ web console opens in a new browser tab or window.
5. Log in to the RabbitMQ web console with your broker administrator sign-in credentials.
6. In the RabbitMQ web console, at the top of the page, choose **Admin**.
7. On the **Admin** page, in the right navigation pane, choose **Policies**.
8. On the **Policies** page, you can see a list of the broker's current **User policies**. Below **User policies**, expand **Add / update a policy**.

> **Note**
>
> By default, Amazon MQ for RabbitMQ clusters are created with an initial broker policy named `ha-all-AWS-OWNED-DO-NOT-DELETE`. Amazon MQ manages this policy to ensure that every queue on the broker is replicated to all three nodes and that queues are synchronized automatically.

9. To create a new broker policy, under **Add / update a policy**, do the following:
   a. For **Name**, enter a name for your policy, for example, **`batch-size-policy`**.
   b. For **Pattern**, enter the regexp pattern **`.*`** so that the policy matches all queues on the broker.
   c. For **Apply to**, choose **Exchanges and queues** from the dropdown list.
   d. For **Priority**, enter an integer greater than all other policies in applied to the vhost. You can apply exactly one set of policy definitions to RabbitMQ queues and exchanges at any given time. RabbitMQ chooses the matching policy with the highest priority value. For more information about policy priorities and how to combine policies, see Policies in the RabbitMQ Server Documentation.
   e. For **Definition**, add the following key-value pairs:
      - **`ha-sync-batch-size`**=*`100`*. Choose **Number** from the dropdown list.

        > **Note**
        >
        > You might need to adjust and calibrate the value of `ha-sync-batch-size` based on the number and size of unsynchronized messages in your queues.

      - **`ha-mode=all`**. Choose **String** from the dropdown list.

> **Important**
>
> The `ha-mode` definition is required for all HA-related policies. Omitting it results in a validation failure.

- **`ha-sync-mode=automatic`**. Choose **String** from the dropdown list.

> **Note**
>
> The `ha-sync-mode` definition is required for all custom policies. If it is omitted, Amazon MQ automatically appends the definition.

f. Choose **Add / update policy**.

10. Confirm that the new policy appears in the list of **User policies**.

To apply an `ha-sync-batch-size` policy using the RabbitMQ management API

1. Sign in to the Amazon MQ console.
2. In the left navigation pane, choose **Brokers**.
3. From the list of brokers, choose the name of the broker to which you want to apply the new policy.
4. On the broker's page, in the **Connections** section, note the **RabbitMQ web console** URL. This is the broker endpoint that you use in an HTTP request.
5. Open a new terminal or command line window of your choice.
6. To create a new broker policy, enter the following `curl` command. This command assumes a queue on the default `/` vhost, which is encoded as `%2F`.

> **Note**
>
> Replace *`username`* and *`password`* with your broker administrator sign-in credentials. You might need to adjust and calibrate the value of `ha-sync-batch-size` (*`100`*) based on the number and size of unsynchronized messages in your queues. Replace the broker endpoint with the URL that you noted previously.

```
curl -i -u username:password -H "content-type:application/json" -XPUT \
-d '{"pattern":".*", "priority":1, "definition":{"ha-sync-batch-size":100, "ha-
mode":"all", "ha-sync-mode":"automatic"}}' \
https://b-589c045f-f8ln-4ab0-a89c-co62e1c32ef8.mq.us-west-2.amazonaws.com/api/
policies/%2F/batch-size-policy
```

7. To confirm that the new policy is added to your broker's user policies, enter the following `curl` command to list all broker policies.

```
curl -i -u username:password https://b-589c045f-f8ln-4ab0-a89c-co62e1c32ef8.mq.us-
west-2.amazonaws.com/api/policies
```

Step 2: Restart the queue sync

After applying a new `ha-sync-batch-size` policy to your broker, restart the queue sync.

To restart the queue sync using the RabbitMQ web console

> **Note**
>
> To open the RabbitMQ web console, see the previous instructions in Step 1 of this tutorial.

1. In the RabbitMQ web console, at the top of the page, choose **Queues**.
2. On the **Queues** page, under **All queues**, locate your paused queue. In the **Features** column, your queue should list the name of the new policy that you created (for example, `batch-size-policy`).
3. To restart the synchronization process with a reduced batch size, choose **Restart sync**.

> **Note**
>
> If synchronization pauses and doesn't finish successfully, try reducing the `ha-sync-batch-size` value and restarting the queue sync again.

Next steps

- Once your queue synchronizes successfully, you can monitor the amount of memory that your RabbitMQ nodes use by viewing the Amazon CloudWatch metric `RabbitMQMemUsed`. You can also view the `RabbitMQMemLimit` metric to monitor a node's memory limit. For more information, see Accessing CloudWatch metrics for Amazon MQ and Logging and monitoring Amazon MQ for RabbitMQ brokers.
- To prevent paused queue synchronization, we recommend keeping queues short and processing messages. For workloads with larger message sizes, we also recommend upgrading your broker instance type to a larger instance size with more memory. For more information about broker instance types and editing broker preferences, see Amazon MQ for RabbitMQ instance types and Editing broker preferences.
- When you create a new Amazon MQ for RabbitMQ broker, Amazon MQ applies a set of default policies and virtual host limits to optimize broker performance. If your broker does not have the recommended default policies and limits, we recommend creating them yourself. For more information about creating default policies and vhost limits, see the section called "Broker defaults".

Related resources

- UpdateBrokerInput – Use this broker property to update a broker instance type using the Amazon MQ API.
- Parameters and Policies (RabbitMQ Server Documentation) – Learn more about RabbitMQ parameters and policies on the RabbitMQ website.
- RabbitMQ Management HTTP API – Learn more about the RabbitMQ management API.

Amazon MQ for RabbitMQ best practices

Use this as a reference to quickly find recommendations for maximizing performance and minimizing throughput costs when working with RabbitMQ brokers on Amazon MQ.

> **Important**
>
> Amazon MQ does not support quorum queues. Enabling the quorum queue feature flag and creating quorum queues will result in data loss.

> **Important**
>
> Currently, Amazon MQ does not support streams, or using structured logging in JSON, introduced in RabbitMQ 3.9.x.

> **Important**
>
> Amazon MQ for RabbitMQ does not support the username "guest", and will delete the default guest account when you create a new broker. Amazon MQ will also periodically delete any customer created account called "guest".

Topics

Enable lazy queues

If you are working with very long queues that process large volumes of messages, enabling *lazy queues* can improve your broker's overall performance.

RabbitMQ's default behavior is to cache messages in memory and to move them to disk only when the broker needs more available memory. This process of moving messages from memory to disk can take time and stops the queue from processing messages. Enabling lazy queues can have a significant impact on speeding up the process of moving messages to disk as lazy queues store messages to disk as soon as possible, resulting in fewer messages cached in memory.

You can enable lazy queues by setting the `queue.declare` arguments at the time of declaration, or by configuring a policy via the RabbitMQ management console. The following example demonstrates declaring a lazy queue using the RabbitMQ Java client library.

```
Map<String, Object> args = new HashMap<String, Object>();
args.put("x-queue-mode", "lazy");
channel.queueDeclare("myqueue", false, false, false, args);
```

> **Note**
>
> Enabling lazy queues can increase disk I/O operations.

Use persistent and durable queues

Persistent messages can help prevent data loss in situations where a broker crashes or restarts. Persistent messages are written to disk as soon as they arrive. Unlike lazy queues, however, persistent messages are cached both in memory and in disk unless more memory is needed by the broker. In cases where more memory is needed, messages are removed from memory by the RabbitMQ broker mechanism that manages storing messages to disk, commonly referred to as the *persistence layer*.

To enable message persistence, you can declare your queues as `durable` and set message delivery mode to `persistent`. The following example demonstrates using the RabbitMQ Java client library to declare a durable queue.

```
boolean durable = true;
channel.queueDeclare("my_queue", durable, false, false, null);
```

Once you have configured your queue as durable, you can send a persistent message to your queue by setting `MessageProperties` to `PERSISTENT_TEXT_PLAIN` as shown in the following example.

```
import com.rabbitmq.client.MessageProperties;

channel.basicPublish("", "my_queue",
            MessageProperties.PERSISTENT_TEXT_PLAIN,
            message.getBytes());
```

Keep queues short

In cluster deployments, queues with a large number of messages can lead to resource overutilization. When a broker is overutilized, rebooting an Amazon MQ for RabbitMQ broker can cause further degradation of performance. If rebooted, overutilized brokers might become unresponsive in the `REBOOT_IN_PROGRESS` state.

During maintenance windows, Amazon MQ performs all maintenance work one node at a time to ensure that the broker remains operational. As a result, queues might need to synchronize as each node resumes operation. During synchronization, messages that need to be replicated to mirrors are loaded into memory from the corresponding Amazon Elastic Block Store (Amazon EBS) volume to be processed in batches. Processing messages in batches lets queues synchronize faster.

If queues are kept short and messages are small, the queues successfully synchronize and resume operation as expected. However, if the amount of data in a batch approaches the node's memory limit, the node raises a high memory alarm, pausing the queue sync. You can confirm memory usage by comparing the `RabbitMemUsed` and `RabbitMqMemLimit` broker node metrics in CloudWatch. Synchronization can't complete until messages are consumed or deleted, or the number of messages in the batch is reduced.

If queue synchronization is paused for a cluster deployment, we recommend consuming or deleting messages to lower the number of messages in queues. Once queue depth is reduced and queue sync completes, the broker status will change to `RUNNING`. To resolve a paused queue sync, you can also apply a policy to reduce the queue synchronization batch-size.

> **Warning**
>
> Do not reboot a broker that is running high on resources.
> If you reboot a broker when queue synchronization is paused, the broker will reinitiate the synchronization process, which can further degrade broker resources as messages are transferred from storage to node memory, and result in the broker becoming unresponsive in the `REBOOT_IN_PROGRESS` state.

Configure acknowledgement and confirmation

When a client application sends confirmation of delivery and consumption of messages back to the broker, it is known as *consumer acknowledgment*. Similarly, the process of sending confirmation to

a publisher is known as *publisher confirm*. Both acknowledgement and confirmation are essential to ensuring data safety when working with RabbitMQ brokers.

Consumer delivery acknowledgement is typically configured on the client application. When working with AMQP 0-9-1, acknowledgement can be enabled by configuring the `basic.consume` or when a message is fetched using the `basic.code` method.

Typically, delivery acknowledgement is enabled in a channel. For example, when working with the RabbitMQ Java client library, you can use the `Channel#basicAck` to set up a simple `basic.ack` positive acknowledgement as shown in the following example.

```
// this example assumes an existing channel instance

boolean autoAck = false;
channel.basicConsume(queueName, autoAck, "a-consumer-tag",
     new DefaultConsumer(channel) {
         @Override
         public void handleDelivery(String consumerTag,
                                    Envelope envelope,
                                    AMQP.BasicProperties properties,
                                    byte[] body)
             throws IOException
         {
             long deliveryTag = envelope.getDeliveryTag();
             // positively acknowledge a single delivery, the message will
             // be discarded
             channel.basicAck(deliveryTag, false);
         }
     });
```

> **Note**
>
> Unacknowledged messages must be cached in memory. You can limit the number of messages that a consumer pre-fetches by configuring pre-fetch settings for a client application.

Configure pre-fetching

You can use the RabbitMQ pre-fetch value to optimize how your consumers consume messages. RabbitMQ implements the channel pre-fetch mechanism provided by AMQP 0-9-1 by applying the

pre-fetch count to consumers as opposed to channels. The pre-fetch value is used to specify how many messages are being sent to the consumer at any given time. By default, RabbitMQ sets an unlimited buffer size for client applications.

There are a variety of factors to consider when setting a pre-fetch count for your RabbitMQ consumers. First, consider your consumers' environment and configuration. Because consumers need to keep all messages in memory as they are being processed, a high pre-fetch value can have a negative impact on your consumers' performance, and in some cases, can result in a consumer potentially crashing all together. Similarly, the RabbitMQ broker itself keeps all messages that it sends cached in memory until it recieves consumer acknowledgement. A high pre-fetch value can cause your RabbitMQ server to run out of memory quickly if automatic acknowledgement is not configured for consumers, and if consumers take a relatively long time to process messages.

With the above considerations in mind, we recommend always setting a pre-fetch value in order to prevent situations where a RabbitMQ broker or its consumers run out of memory due to a large number number of unprocessed, or unacknowledged messages. If you need to optimize your brokers to process large volumes of messages, you can test your brokers and consumers using a range of pre-fetch counts to determine the value at which point network overhead becomes largely insignificant compared to the time it takes a consumer to process messages.

Note

- If your client applications have configured to automatically acknowledge delivery of messages to consumers, setting a pre-fetch value will have no effect.
- All pre-fetched messages are removed from the queue.

The following example desmonstrate setting a pre-fetch value of `10` for a single consumer using the RabbitMQ Java client library.

```
ConnectionFactory factory = new ConnectionFactory();

Connection connection = factory.newConnection();
Channel channel = connection.createChannel();

channel.basicQos(10, false);

QueueingConsumer consumer = new QueueingConsumer(channel);
```

```
channel.basicConsume("my_queue", false, consumer);
```

> **Note**
>
> In the RabbitMQ Java client library, the default value for the `global` flag is set to `false`, so the above example can be written simply as `channel.basicQos(10)`.

Configure Celery

Python Celery sends a lot of unnecessary messages that can make finding and processing the useful information harder. To reduce the noise and make processing easier, enter the following command:

```
celery -A app_name worker --without-heartbeat --without-gossip --without-mingle
```

Automatically recover from network failures

We recommend always enabling automatic network recovery to prevent significant downtime in cases where client connections to RabbitMQ nodes fail. The RabbitMQ Java client library supports automatic network recovery by default, beginning with version `4.0.0`.

Automatic connection recovery is triggered if an unhandled exception is thrown in the connection's I/O loop, if a socket read operation timeout is detected, or if the server misses a heartbeat.

In cases where the initial connection between a client and a RabbitMQ node fails, automatic recovery will not be triggered. We recommend writing your application code to account for initial connection failures by retrying the connection. The following example demonstrates retrying initial network failures using the RabbitMQ Java client library.

```
ConnectionFactory factory = new ConnectionFactory();
// enable automatic recovery if using RabbitMQ Java client library prior to version
 4.0.0.
factory.setAutomaticRecoveryEnabled(true);
// configure various connection settings

try {
  Connection conn = factory.newConnection();
```

```
} catch (java.net.ConnectException e) {
  Thread.sleep(5000);
  // apply retry logic
}
```

> **Note**
>
> If an application closes a connection by using the `Connection.Close` method, automatic network recovery will not be enabled or triggered.

Enable Classic Queue v2 for your RabbitMQ broker

We recommend enabling Classic Queue v2 (CQv2) for broker engine versions 3.10 or above for performance improvements including:

- Flatten memory usage variation
- Decrease memory usage for most workloads
- Improve consumer delivery
- Increase throughput for workloads where consumers keep up with producers

To use CQv2, you must first enable the `classic_mirrored_queue_version` feature flag. For more information on feature flags, see How to enable feature flags.

To migrate from CQv1 to CQv2, you must create a new queue policy or edit an existing queue policy with the `queue-version` policy key definition set to 2. For more information on applying policies, see Policies. For more information on enabling CQv2 with a queue policy, see Classic Queues in the RabbitMQ documentation.

Migrating from CQv1 to CQv2 can be memory intensive for queues under a high load. We recommend following our other best performance practices before starting the migration.

If you are using a queue policy, deleting the queue policy will downgrade CQv2 queues back to CQv1. We do not recommend downgrading CQv2 queues to CQv1 because RabbitMQ will convert the on-disk representation of the queue. This can be memory intensive and time-consuming for queues with high depth.

Quotas in Amazon MQ for RabbitMQ

This topic lists quotas within Amazon MQ. Many of the following quotas can be changed for specific AWS accounts. To request an increase for a limit, see AWS Service Quotas in the *Amazon Web Services General Reference*. Updated limits will not be visible even after the limit increase has been applied. For more information on viewing current connection limits in Amazon CloudWatch, see Monitoring Amazon MQ brokers using Amazon CloudWatch.

Topics

Brokers

The following table lists quotas related to Amazon MQ for RabbitMQ brokers.

Limit	Description
Broker name	• Must be unique in the broker region and your AWS account. • Must be 1-50 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~).
Number of brokers, per region	50
Security groups per broker	5

Limit	Description
ActiveMQ destinations (queues, and topics) monitored in CloudWatch	CloudWatch monitors only the first 1000 destinations.
RabbitMQ destinations (queues) monitored in CloudWatch	CloudWatch monitors only the first 500 destinations, ordered by number of consumers .
Tags per broker	50

Data Storage

The following table lists quotas related to Amazon MQ for RabbitMQ data storage.

Limit	Description
Storage capacity per smaller broker	20 GB for `mq.*.micro` instance type brokers. For more information regarding A mazon MQ instance types, see Broker instance types.
Storage capacity per larger broker	200 GB for `mq.*.*large` instance type brokers. For more information regarding A mazon MQ instance types, see Broker instance types.

API Throttling

The following throttling quotas are aggregated per AWS account, *across all Amazon MQ APIs* to maintain service bandwidth. For more information about Amazon MQ APIs, see the *Amazon MQ REST API Reference*.

Important

These quotas don't apply to Amazon MQ for ActiveMQ or Amazon MQ for RabbitMQ broker messaging APIs. For example, Amazon MQ doesn't throttle the sending or receiving of messages.

API burst limit	API rate limit
100	15

Security in Amazon MQ

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The shared responsibility model describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the AWS Compliance Programs. To learn about the compliance programs that apply to Amazon MQ, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon MQ. The following topics show you how to configure Amazon MQ to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon MQ resources.

Topics

Data protection in Amazon MQ

The AWS shared responsibility model applies to data protection in Amazon MQ. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud.

You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the Data Privacy FAQ. For information about data protection in Europe, see the AWS Shared Responsibility Model and GDPR blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-2.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon MQ or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

For both Amazon MQ for ActiveMQ and Amazon MQ for RabbitMQ brokers, do not use any personally identifiable information (PII) or other confidential or sensitive information for broker names or usernames when creating resources via the broker web console, or the Amazon MQ API. Broker names and usernames are accessible to other AWS services, including CloudWatch Logs. Broker usernames are not intended to be used for private or sensitive data.

Encryption

User data stored in Amazon MQ is encrypted at rest. Amazon MQ encryption at rest provides enhanced security by encrypting your data using encryption keys stored in the AWS Key Management Service (KMS). This service helps reduce the operational burden and complexity involved in protecting sensitive data. With encryption at rest, you can build security-sensitive applications that meet encryption compliance and regulatory requirements.

All connections between Amazon MQ brokers use Transport layer Security (TLS) to provide encryption in transit.

Amazon MQ encrypts messages at rest and in transit using encryption keys that it manages and stores securely. For more information, see the *AWS Encryption SDK Developer Guide*.

Encryption at rest

Amazon MQ integrates with AWS Key Management Service (KMS) to offer transparent server-side encryption. Amazon MQ always encrypts your data at rest.

When you create an Amazon MQ for ActiveMQ broker or an Amazon MQ for RabbitMQ broker, you can specify the AWS KMS key that you want Amazon MQ to use to encrypt your data at rest. If you do not specify a KMS key, Amazon MQ creates an AWS owned KMS key for you and uses it on your behalf. Amazon MQ currently supports symmetric KMS keys. For more information about KMS keys, see AWS KMS keys.

When creating a broker, you can configure what Amazon MQ uses for your encryption key by selecting one of the following.

- **Amazon MQ owned KMS key (default)** — The key is owned and managed by Amazon MQ and is not in your account.
- **AWS managed KMS key** — The AWS managed KMS key (`aws/mq`) is a KMS key in your account that is created, managed, and used on your behalf by Amazon MQ.
- **Select existing customer managed KMS key** — Customer managed KMS keys are created and managed by you in AWS Key Management Service (KMS).

Important

- Revoking a grant cannot be undone. Instead, we suggest deleting the broker if you need to revoke access rights.
- For **Amazon MQ for ActiveMQ** brokers that use Amazon Elastic File System (EFS) to store message data, if you revoke the grant that gives Amazon EFS permission to use the KMS keys in your account, it will not take place immediately.
- For **Amazon MQ for RabbitMQ** and **Amazon MQ for ActiveMQ** brokers that use EBS to store message data, if you disable, schedule for deletion, or revoke the grant that gives Amazon EBS permission to use the KMS keys in your account, Amazon MQ cannot maintain your broker, and it may change to a degraded state.
- If you have deactivated the key or scheduled the key to be deleted, you can reactivate the key or cancel key deletion and keep your broker maintained.
- Deactivating a key or revoking a grant will not take place immediately.

When creating a single instance broker with a KMS key for RabbitMQ, you will see two `CreateGrant` events logged in AWS CloudTrail. The first event is Amazon MQ creating a grant for the KMS key. The second event is EBS creating a grant for EBS to use.

CreateGrant AWS CloudTrail log entry: single instance broker

mq_grant

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AKIAIOSFODNN7EXAMPLE",
        "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
        "accountId": "111122223333",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AKIAIOSFODNN7EXAMPLE",
                "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
                "accountId": "111122223333",
```

```
                "userName": "AmazonMqConsole"
            },
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2023-02-23T18:59:10Z",
                "mfaAuthenticated": "false"
            }
        },
        "invokedBy": "mq.amazonaws.com"
    },
    "eventTime": "2018-06-28T22:23:46Z",
    "eventSource": "amazonmq.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "PostmanRuntime/7.1.5",
    "requestParameters": {
        "granteePrincipal": "mq.amazonaws.com",
        "keyId": "arn:aws:kms:us-east-1:316438333700:key/bdbe42ae-f825-4e78-
a8a1-828d411c4be2",
        "retiringPrincipal": "mq.amazonaws.com",
        "operations": [
            "CreateGrant",
            "Decrypt",
            "GenerateDataKeyWithoutPlaintext",
            "ReEncryptFrom",
            "ReEncryptTo",
            "DescribeKey"
        ]
    },
    "responseElements": {
        "grantId":
 "0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",

    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": false,
    "resources": [
        {
           "accountId": "111122223333",
            "type": "AWS::KMS::Key",
```

```
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management",
    "sessionCredentialFromConsole": "true"
}
```

EBS grant creation

You will see one event for EBS grant creation.

```
                                    {
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "mq.amazonaws.com"
    },
    "eventTime": "2023-02-23T19:09:40Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "mq.amazonaws.com",
    "userAgent": "ExampleDesktop/1.0 (V1; OS)",
    "requestParameters": {
        "granteePrincipal": "mq.amazonaws.com",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
        "constraints": {
            "encryptionContextSubset": {
                "aws:ebs:id": "vol-0b670f00f7d5417c0"
            }
        },
        "operations": [
            "Decrypt"
        ],
        "retiringPrincipal": "ec2.us-east-1.amazonaws.com"
    },
    "responseElements": {
```

```
        "grantId":
 "0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
    },
    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": false,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "sharedEventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventCategory": "Management"
}
```

When creating a cluster deployment with a KMS key for RabbitMQ, you will see five `CreateGrant` events logged in AWS CloudTrail. The first two events are grant creations for Amazon MQ. The next three events are grants created by EBS for EBS to use.

CreateGrant AWS CloudTrail log entry: cluster deployment

mq_grant

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AKIAIOSFODNN7EXAMPLE",
        "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
        "accountId": "111122223333",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
```

```
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AKIAIOSFODNN7EXAMPLE",
                "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
                "accountId": "111122223333",
                "userName": "AmazonMqConsole"
            },
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2023-02-23T18:59:10Z",
                "mfaAuthenticated": "false"
            }
        },
        "invokedBy": "mq.amazonaws.com"
    },
    "eventTime": "2018-06-28T22:23:46Z",
    "eventSource": "amazonmq.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "PostmanRuntime/7.1.5",
    "requestParameters": {
        "granteePrincipal": "mq.amazonaws.com",
        "keyId": "arn:aws:kms:us-east-1:316438333700:key/bdbe42ae-f825-4e78-
a8a1-828d411c4be2",
        "retiringPrincipal": "mq.amazonaws.com",
        "operations": [
            "CreateGrant",
            "Encrypt",
            "Decrypt",
            "ReEncryptFrom",
            "ReEncryptTo",
            "GenerateDataKey",
            "GenerateDataKeyWithoutPlaintext",
            "DescribeKey"
        ]
    },
    "responseElements": {
        "grantId":
 "0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",

    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
```

```
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": false,
    "resources": [
        {
           "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management",
    "sessionCredentialFromConsole": "true"
}
```

mq_rabbit_grant

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AssumedRole",
        "principalId": "AKIAIOSFODNN7EXAMPLE",
        "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
        "accountId": "111122223333",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
            "sessionIssuer": {
                "type": "Role",
                "principalId": "AKIAIOSFODNN7EXAMPLE",
                "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
                "accountId": "111122223333",
                "userName": "AmazonMqConsole"
            },
            "webIdFederationData": {},
            "attributes": {
                "creationDate": "2023-02-23T18:59:10Z",
                "mfaAuthenticated": "false"
            }
        },
```

```
        "invokedBy": "mq.amazonaws.com"
    },
    "eventTime": "2018-06-28T22:23:46Z",
    "eventSource": "amazonmq.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-west-2",
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "PostmanRuntime/7.1.5",
    "requestParameters": {
        "granteePrincipal": "mq.amazonaws.com",
        "retiringPrincipal": "mq.amazonaws.com",
        "operations": [
            "DescribeKey"
        ],
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
    },
    "responseElements": {
        "grantId":
 "0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",

    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": false,
    "resources": [
        {
           "accountId": "111122223333",
            "type": "AWS::KMS::Key",
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "eventCategory": "Management",
    "sessionCredentialFromConsole": "true"
}
```

EBS grant creation

You will see three events for EBS grant creation.


```
                                    {
    "eventVersion": "1.08",
    "userIdentity": {
        "type": "AWSService",
        "invokedBy": "mq.amazonaws.com"
    },
    "eventTime": "2023-02-23T19:09:40Z",
    "eventSource": "kms.amazonaws.com",
    "eventName": "CreateGrant",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "mq.amazonaws.com",
    "userAgent": "ExampleDesktop/1.0 (V1; OS)",
    "requestParameters": {
        "granteePrincipal": "mq.amazonaws.com",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
        "constraints": {
            "encryptionContextSubset": {
                "aws:ebs:id": "vol-0b670f00f7d5417c0"
            }
        },
        "operations": [
            "Decrypt"
        ],
        "retiringPrincipal": "ec2.us-east-1.amazonaws.com"
    },
    "responseElements": {
        "grantId":
 "0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE",
        "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
    },
    "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "readOnly": false,
    "resources": [
        {
            "accountId": "111122223333",
            "type": "AWS::KMS::Key",
```

```
            "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
        }
    ],
    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "111122223333",
    "sharedEventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
    "eventCategory": "Management"
}
```

For more information about KMS keys, see AWS KMS keys in the *AWS Key Management Service Developer Guide*.

Encryption in transit

Amazon MQ for ActiveMQ: Amazon MQ for ActiveMQ requires strong Transport Layer Security (TLS) and encrypts data in transit between the brokers of your Amazon MQ deployment. All data that passes between Amazon MQ brokers is encrypted using strong Transport Layer Security (TLS). This is true for all available protocols.

Amazon MQ for RabbitMQ: Amazon MQ for RabbitMQ requires strong Transport Layer Security (TLS) encryption for all client connections. RabbitMQ cluster replication traffic only transits your broker's VPC and all network traffic between AWS data centers is transparently encrypted at the physical layer. Amazon MQ for RabbitMQ clustered brokers currently do not support Inter-node encryption for cluster replication. To learn more about data-in-transit, see Encrypting Data-at-Rest and -in-Transit.

Amazon MQ for ActiveMQ protocols

You can access your ActiveMQ brokers using the following protocols with TLS enabled:

- AMQP
- MQTT
- MQTT over WebSocket
- OpenWire
- STOMP

- STOMP over WebSocket

Supported TLS Cipher Suites for ActiveMQ

ActiveMQ on Amazon MQ supports the following cipher suites:

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA

Amazon MQ for RabbitMQ protocols

You can access your RabbitMQ brokers using the following protocols with TLS enabled:

- AMQP (0-9-1)

Supported TLS Cipher Suites for RabbitMQ

RabbitMQ on Amazon MQ supports the following cipher suites:

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256

Identity and access Management for Amazon MQ

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon MQ resources. IAM is an AWS service that you can use with no additional charge.

Topics

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Amazon MQ.

Service user – If you use the Amazon MQ service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon MQ features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon MQ, see Troubleshooting Amazon MQ identity and access.

Service administrator – If you're in charge of Amazon MQ resources at your company, you probably have full access to Amazon MQ. It's your job to determine which Amazon MQ features and resources your service users should access. You must then submit requests to your IAM

administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon MQ, see How Amazon MQ works with IAM.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon MQ. To view example Amazon MQ identity-based policies that you can use in IAM, see Amazon MQ Identity-based policy examples.

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see How to sign in to your AWS account in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see Signing AWS API requests in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see Multi-factor authentication in the *AWS IAM Identity Center User Guide* and Using multi-factor authentication (MFA) in AWS in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and

is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see Tasks that require root user credentials in the *IAM User Guide*.

Users and groups

An *IAM user* is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see Rotate access keys regularly for use cases that require long-term credentials in the *IAM User Guide*.

An *IAM group* is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see When to create an IAM user (instead of a role) in the *IAM User Guide*.

IAM roles

An *IAM role* is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by switching roles. You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see Using IAM roles in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see Creating a role for a third-party Identity Provider

in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see Permission sets in the *AWS IAM Identity Center User Guide*.

- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
  - **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.
  - **Service role** – A service role is an IAM role that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see Creating a role to delegate permissions to an AWS service in the *IAM User Guide*.
  - **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role

to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Using an IAM role to grant permissions to applications running on Amazon EC2 instances in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see When to create an IAM role (instead of a user) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see Overview of JSON policies in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Creating IAM policies in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone

policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see Choosing between managed policies and inline policies in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must specify a principal in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access Control Lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see Access control list (ACL) overview in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see Permissions boundaries for IAM entities in the *IAM User Guide*.

- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see How SCPs work in the *AWS Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see Session policies in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

How Amazon MQ works with IAM

Before you use IAM to manage access to Amazon MQ, you should understand what IAM features are available to use with Amazon MQ. To get a high-level view of how Amazon MQ and other AWS services work with IAM, see AWS Services That Work with IAM in the *IAM User Guide*.

Amazon MQ uses IAM for creating, updating, and deleting operations, but native ActiveMQ authentication for brokers. For more information, see Integrating ActiveMQ brokers with LDAP.

Topics

Amazon MQ identity-based policies

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. Amazon MQ supports specific actions, resources, and condition keys. To learn about all of the elements that you use in a JSON policy, see IAM JSON Policy Elements Reference in the *IAM User Guide*.

Actions

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy actions in Amazon MQ use the following prefix before the action: `mq:`. For example, to grant someone permission to run an Amazon MQ instance with the Amazon MQ `CreateBroker` API operation, you include the `mq:CreateBroker` action in their policy. Policy statements must include either an `Action` or `NotAction` element. Amazon MQ defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas as follows:

```
"Action": [
      "mq:action1",
      "mq:action2"
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Describe`, include the following action:

```
"Action": "mq:Describe*"
```

To see a list of Amazon MQ actions, see Actions Defined by Amazon MQ in the *IAM User Guide*.

Resources

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Resource` JSON policy element specifies the object or objects to which the action applies. Statements must include either a `Resource` or a `NotResource` element. As a best practice, specify a resource using its Amazon Resource Name (ARN). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

In the Amazon MQ, the primary AWS resources are an Amazon MQ message broker and its configuration. Amazon MQ brokers and configurations each have unique Amazon Resource Names (ARNs) associated with them, as shown in the following table.

Resource Types	ARN	Condition Keys
brokers	`arn:aws:mq:us-east-1:123456789012:broker:${brokerName}:${brokerId}`	aws:ResourceTag/${TagKey}
configurations	`arn:${Partition}:mq:${Region}:${Account}:configuration:${configuration-id}`	aws:ResourceTag/${TagKey}

For more information about the format of ARNs, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

For example, to specify the broker named `MyBroker` with brokerId `b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9` in your statement, use the following ARN:

```
"Resource": "arn:aws:mq:us-
east-1:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9"
```

To specify all brokers and configurations that belong to a specific account, use the wildcard (*):

```
"Resource": "arn:aws:mq:us-east-1:123456789012:*"
```

Some Amazon MQ actions, such as those for creating resources, cannot be performed on a specific resource. In those cases, you must use the wildcard (*).

```
"Resource": "*"
```

The API action `CreateTags` requires both a broker and a configuration. To specify multiple resources in a single statement, separate the ARNs with commas.

```
"Resource": [
      "resource1",
      "resource2"
```

To see a list of Amazon MQ resource types and their ARNs, see Resources Defined by Amazon MQ in the *IAM User Guide*. To learn with which actions you can specify the ARN of each resource, see Actions Defined by Amazon MQ.

Condition keys

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or `Condition` *block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use condition operators, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical `AND` operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical `OR` operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see AWS global condition context keys in the *IAM User Guide*.

Amazon MQ does not define any service-specific condition keys, but supports using some global condition keys. To see a list of Amazon MQ condition keys, see the table below or Condition Keys for Amazon MQ in the *IAM User Guide*. To learn with which actions and resources you can use a condition key, see Actions Defined by Amazon MQ.

Condition Keys	Description	Type
aws:RequestTag/${TagKey}	Filters actions based on the tags that are passed in the request.	String
aws:ResourceTag/${TagKey}	Filters actions based on the tags associated with the resource.	String
aws:TagKeys	Filters actions based on the tag keys that are passed in the request.	String

Examples

To view examples of Amazon MQ identity-based policies, see Amazon MQ Identity-based policy examples.

Amazon MQ Resource-based policies

Currently, Amazon MQ doesn't support IAM authentication using resource-based permissions or resource-based policies.

Authorization based on Amazon MQ tags

You can attach tags to Amazon MQ resources or pass tags in a request to Amazon MQ. To control access based on tags, you provide tag information in the condition element of a policy using the `mq:ResourceTag/`*`key-name`*, `aws:RequestTag/`*`key-name`*, or `aws:TagKeys` condition keys.

Amazon MQ supports policies based on tags. For instance, you could deny access to Amazon MQ resources that include a tag with the key `environment` and the value `production`:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": [
                "mq:DeleteBroker",
                "mq:RebootBroker",
                "mq:DeleteTags"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "aws:ResourceTag/environment": "production"
                }
            }
        }
    ]
}
```

This policy will `Deny` the ability to delete or reboot an Amazon MQ broker that includes the tag `environment/production`.

For more information on tagging, see:

- Tagging resources
- Controlling Access Using IAM Tags

Amazon MQ IAM roles

An IAM role is an entity within your AWS account that has specific permissions.

Using Temporary Credentials with Amazon MQ

You can use temporary credentials to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as AssumeRole or GetFederationToken.

Amazon MQ supports using temporary credentials.

Service roles

This feature allows a service to assume a service role on your behalf. This role allows the service to access resources in other services to complete an action on your behalf. Service roles appear in your IAM account and are owned by the account. This means that an IAM administrator can change the permissions for this role. However, doing so might break the functionality of the service.

Amazon MQ supports service roles.

Amazon MQ Identity-based policy examples

By default, users and roles don't have permission to create or modify Amazon MQ resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the specified resources they need. The administrator must then attach those policies to the IAM users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see Creating Policies on the JSON Tab in the *IAM User Guide*.

Topics

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon MQ resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see AWS managed policies or AWS managed policies for job functions in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see Policies and permissions in IAM in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see IAM JSON policy elements: Condition in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see IAM Access Analyzer policy validation in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see Configuring MFA-protected API access in the *IAM User Guide*.

For more information about best practices in IAM, see Security best practices in IAM in the *IAM User Guide*.

Using the Amazon MQ console

To access the Amazon MQ console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon MQ resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (IAM users or roles) with that policy.

To ensure that those entities can still use the Amazon MQ console, also attach the following AWS managed policy to the entities. For more information, see Adding Permissions to a User in the *IAM User Guide*:

```
AmazonMQReadOnlyAccess
```

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that you're trying to perform.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
        }
    ]
}
```

API authentication and authorization for Amazon MQ

Amazon MQ uses standard AWS request signing for API authentication. For more information, see Signing AWS API Requests in the *AWS General Reference*.

> **Note**
>
> Currently, Amazon MQ doesn't support IAM authentication using resource-based permissions or resource-based policies.

To authorize AWS users to work with brokers, configurations, and users, you must edit your IAM policy permissions.

Topics

IAM Permissions Required to Create an Amazon MQ Broker

To create a broker, you must either use the `AmazonMQFullAccess` IAM policy or include the following EC2 permissions in your IAM policy.

The following custom policy is comprised of two statements (one conditional) which grant permissions to manipulate the resources which Amazon MQ requires to create an ActiveMQ broker.

> **Important**
>
> - The `ec2:CreateNetworkInterface` action is required to allow Amazon MQ to create an elastic network interface (ENI) in your account on your behalf.
> - The `ec2:CreateNetworkInterfacePermission` action authorizes Amazon MQ to attach the ENI to an ActiveMQ broker.
> - The `ec2:AuthorizedService` condition key ensures that ENI permissions can be granted only to Amazon MQ service accounts.

```
{
    "Version": "2012-10-17",
    "Statement": [{
        "Action": [
            "mq:*",
            "ec2:CreateNetworkInterface",
            "ec2:DeleteNetworkInterface",
            "ec2:DetachNetworkInterface",
            "ec2:DescribeInternetGateways",
            "ec2:DescribeNetworkInterfaces",
            "ec2:DescribeRouteTables",
            "ec2:DescribeSecurityGroups",
            "ec2:DescribeSubnets",
            "ec2:DescribeVpcs"
        ],
        "Effect": "Allow",
        "Resource": "*"
    },{
        "Action": [
            "ec2:CreateNetworkInterfacePermission",
            "ec2:DeleteNetworkInterfacePermission",
            "ec2:DescribeNetworkInterfacePermissions"
        ],
        "Effect": "Allow",
        "Resource": "*",
        "Condition": {
            "StringEquals": {
                "ec2:AuthorizedService": "mq.amazonaws.com"
            }
        }
    }]
}
```

For more information, see Step 2: create a user and get your AWS credentials and Never Modify or Delete the Amazon MQ Elastic Network Interface.

Amazon MQ REST API permissions reference

The following table lists Amazon MQ REST APIs and the corresponding IAM permissions.

Amazon MQ REST APIs and Required Permissions

Amazon MQ REST APIs	Required Permissions
CreateBroker	mq:CreateBroker
CreateConfiguration	mq:CreateConfiguration
CreateTags	mq:CreateTags
CreateUser	mq:CreateUser
DeleteBroker	mq:DeleteBroker
DeleteUser	mq:DeleteUser
DescribeBroker	mq:DescribeBroker
DescribeConfiguration	mq:DescribeConfiguration
DescribeConfigurationRevision	mq:DescribeConfigurationRevision
DescribeUser	mq:DescribeUser
ListBrokers	mq:ListBrokers
ListConfigurationRevisions	mq:ListConfigurationRevisions
ListConfigurations	mq:ListConfigurations
ListTags	mq:ListTags
ListUsers	mq:ListUsers
RebootBroker	mq:RebootBroker
UpdateBroker	mq:UpdateBroker
UpdateConfiguration	mq:UpdateConfiguration
UpdateUser	mq:UpdateUser

Resource-level permissions for Amazon MQ API actions

The term *resource-level permissions* refers to the ability to specify the resources on which users are allowed to perform actions. Amazon MQ has partial support for resource-level permissions. For certain Amazon MQ actions, you can control when users are allowed to use those actions based on conditions that have to be fulfilled, or specific resources that users are allowed to use.

The following table describes the Amazon MQ API actions that currently support resource-level permissions, as well as the supported resources, resource ARNs, and condition keys for each action.

> **Important**
>
> If an Amazon MQ API action is not listed in this table, then it does not support resource-level permissions. If an Amazon MQ API action does not support resource-level permissions, you can grant users permission to use the action, but you have to specify a * wildcard for the resource element of your policy statement.

API Action	Resource Types (*required)
`CreateConfiguration`	configurations*
`CreateTags`	brokers, configurations
`CreateUser`	brokers*
`DeleteBroker`	brokers*
`DeleteUser`	brokers*
`DescribeBroker`	brokers*
`DescribeConfiguration`	configurations*
`DescribeConfigurationRevision`	configurations*
`DescribeUser`	brokers*

API Action	Resource Types (*required)
ListConfigurationRevisions	configurations*
ListConfigurationRevisions	configurations*
ListTags	brokers, configurations
ListUsers	brokers*
RebootBroker	brokers*
UpdateBroker	brokers*
UpdateConfiguration	configurations*
UpdateUser	brokers*

AWS managed policies for Amazon MQ

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining customer managed policies that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see AWS managed policies in the *IAM User Guide*.

AWS managed policy: AmazonMQServiceRolePolicy

You can't attach `AmazonMQServiceRolePolicy` to your IAM entities. This policy is attached to a service-linked role that allows Amazon MQ to perform actions on your behalf. For more information about this permission policy and the actions it allows Amazon MQ to perform, see the section called "Service-linked role permissions for Amazon MQ".

Amazon MQ updates to AWS managed policies

View details about updates to AWS managed policies for Amazon MQ since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Amazon MQ Document history page.

Change	Description	Date
Amazon MQ started tracking changes	Amazon MQ started tracking changes for its AWS managed policies.	May 5, 2021

Using service-linked roles for Amazon MQ

Amazon MQ uses AWS Identity and Access Management (IAM) service-linked roles. A service-linked role is a unique type of IAM role that is linked directly to Amazon MQ. Service-linked roles are predefined by Amazon MQ and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up Amazon MQ easier because you don't have to manually add the necessary permissions. Amazon MQ defines the permissions of its service-linked roles, and unless defined otherwise, only Amazon MQ can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your Amazon MQ resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see AWS services that work with IAM and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for Amazon MQ

Amazon MQ uses the service-linked role named **AWSServiceRoleForAmazonMQ** – Amazon MQ uses this service-linked role to call AWS services on your behalf.

The AWSServiceRoleForAmazonMQ service-linked role trusts the following services to assume the role:

- `mq.amazonaws.com`

Amazon MQ uses the permission policy `AmazonMQServiceRolePolicy`, which is attached to the AWSServiceRoleForAmazonMQ service-linked role, to complete the following actions on the specified resources:

- Action: `ec2:CreateVpcEndpoint` on the `vpc` resource.
- Action: `ec2:CreateVpcEndpoint` on the `subnet` resource.
- Action: `ec2:CreateVpcEndpoint` on the `security-group` resource.
- Action: `ec2:CreateVpcEndpoint` on the `vpc-endpoint` resource.
- Action: `ec2:DescribeVpcEndpoints` on the `vpc` resource.
- Action: `ec2:DescribeVpcEndpoints` on the `subnet` resource.
- Action: `ec2:CreateTags` on the `vpc-endpoint` resource.
- Action: `logs:PutLogEvents` on the `log-group` resource.
- Action: `logs:DescribeLogStreams` on the `log-group` resource.
- Action: `logs:DescribeLogGroups` on the `log-group` resource.
- Action: `CreateLogStream` on the `log-group` resource.

- Action: `CreateLogGroup` on the `log-group` resource.

When you create an Amazon MQ for RabbitMQ broker, the `AmazonMQServiceRolePolicy` permission policy allows Amazon MQ to perform the following tasks on your behalf.

- Create a Amazon VPC endpoint for the broker using the Amazon VPC, subnet, and security-group you provide. You can use the endpoint created for your broker to connect to the broker via the RabbitMQ management console, the management API, or programatically.
- Create log groups, and publish broker logs to Amazon CloudWatch Logs.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeVpcEndpoints"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:CreateVpcEndpoint"
            ],
            "Resource": [
                "arn:aws:ec2:*:*:vpc/*",
                "arn:aws:ec2:*:*:subnet/*",
                "arn:aws:ec2:*:*:security-group/*"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:CreateVpcEndpoint"
            ],
            "Resource": [
                "arn:aws:ec2:*:*:vpc-endpoint/*"
            ],
```

```
            "Condition": {
                "StringEquals": {
                    "aws:RequestTag/AMQManaged": "true"
                }
            }
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:CreateTags"
            ],
            "Resource": "arn:aws:ec2:*:*:vpc-endpoint/*",
            "Condition": {
                "StringEquals": {
                    "ec2:CreateAction": "CreateVpcEndpoint"
                }
            }
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:DeleteVpcEndpoints"
            ],
            "Resource": "arn:aws:ec2:*:*:vpc-endpoint/*",
            "Condition": {
                "StringEquals": {
                    "ec2:ResourceTag/AMQManaged": "true"
                }
            }
        },
        {
            "Effect": "Allow",
            "Action": [
                "logs:PutLogEvents",
                "logs:DescribeLogStreams",
                "logs:DescribeLogGroups",
                "logs:CreateLogStream",
                "logs:CreateLogGroup"
            ],
            "Resource": [
                "arn:aws:logs:*:*:log-group:/aws/amazonmq/*"
            ]
        }
    ]
```

```
}
```

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see Service-Linked Role Permissions in the *IAM User Guide*.

Creating a service-linked role for Amazon MQ

You don't need to manually create a service-linked role. When you first create a broker, Amazon MQ creates a service-linked role to call AWS services on your behalf. All subsequent brokers that you create will use the same role and no new role is created.

> **Important**
>
> This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. To learn more, see A New Role Appeared in My IAM Account.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account.

You can also use the IAM console to create a service-linked role with the **Amazon MQ** use case. In the AWS CLI or the AWS API, create a service-linked role with the `mq.amazonaws.com` service name. For more information, see Creating a service-linked role in the *IAM User Guide*. If you delete this service-linked role, you can use this same process to create the role again.

Editing a service-linked role for Amazon MQ

Amazon MQ does not allow you to edit the AWSServiceRoleForAmazonMQ service-linked role. However, you can edit the description of the role using IAM. For more information, see Editing a service-linked role in the *IAM User Guide*.

Deleting a service-linked role for Amazon MQ

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don’t have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

> **Note**
>
> If the Amazon MQ service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

To delete Amazon MQ resources used by the AWSServiceRoleForAmazonMQ

- Delete your Amazon MQ brokers using the AWS Management Console, Amazon MQ CLI, or Amazon MQ API. For more information about deleting brokers, see ???.

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the AWSServiceRoleForAmazonMQ service-linked role. For more information, see Deleting a Service-Linked Role in the *IAM User Guide*.

Supported regions for Amazon MQ service-linked roles

Amazon MQ supports using service-linked roles in all of the regions where the service is available. For more information, see AWS Regions and Endpoints.

Region name	Region identity	Support in Amazon MQ
US East (N. Virginia)	us-east-1	Yes
US East (Ohio)	us-east-2	Yes
US West (N. California)	us-west-1	Yes
US West (Oregon)	us-west-2	Yes
Asia Pacific (Mumbai)	ap-south-1	Yes
Asia Pacific (Osaka)	ap-northeast-3	Yes
Asia Pacific (Seoul)	ap-northeast-2	Yes
Asia Pacific (Singapore)	ap-southeast-1	Yes
Asia Pacific (Sydney)	ap-southeast-2	Yes

Region name	Region identity	Support in Amazon MQ
Asia Pacific (Tokyo)	ap-northeast-1	Yes
Canada (Central)	ca-central-1	Yes
Europe (Frankfurt)	eu-central-1	Yes
Europe (Ireland)	eu-west-1	Yes
Europe (London)	eu-west-2	Yes
Europe (Paris)	eu-west-3	Yes
South America (São Paulo)	sa-east-1	Yes
AWS GovCloud (US)	us-gov-west-1	No

Troubleshooting Amazon MQ identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon MQ and IAM.

Topics

I Am Not Authorized to Perform an Action in Amazon MQ

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your sign-in credentials.

The following example error occurs when the `mateojackson` user tries to use the console to view details about a *`widget`* but does not have `mq:`*`GetWidget`* permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
 mq:GetWidget on resource: my-example-widget
```

In this case, Mateo asks his administrator to update his policies to allow him to access the *my-example-widget* resource using the `mq:GetWidget` action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon MQ.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon MQ. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
 iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon MQ resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon MQ supports these features, see How Amazon MQ works with IAM.
- To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the *IAM User Guide*.

- To learn how to provide access to your resources to third-party AWS accounts, see Providing access to AWS accounts owned by third parties in the *IAM User Guide*.
- To learn how to provide access through identity federation, see Providing access to externally authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the *IAM User Guide*.

Compliance validation for Amazon MQ

Third-party auditors assess the security and compliance of Amazon MQ as part of multiple AWS compliance programs. These include SOC, PCI, HIPAA, and others.

To learn whether an AWS service is within the scope of specific compliance programs, see AWS services in Scope by Compliance Program and choose the compliance program that you are interested in. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- Security and Compliance Quick Start Guides – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on AWS that are security and compliance focused.
- Architecting for HIPAA Security and Compliance on Amazon Web Services – This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.

  > **Note**
  >
  > Not all AWS services are HIPAA eligible. For more information, see the HIPAA Eligible Services Reference.

- AWS Compliance Resources – This collection of workbooks and guides might apply to your industry and location.
- AWS Customer Compliance Guides – Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map

the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).

- Evaluating Resources with Rules in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- AWS Security Hub – This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see Security Hub controls reference.
- AWS Audit Manager – This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in Amazon MQ

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

Infrastructure security in Amazon MQ

As a managed service, is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see AWS Cloud Security. To design your AWS environment using the best practices for infrastructure security, see Infrastructure Protection in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.

- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests.

Security best practices for Amazon MQ

The following design patterns can improve the security of your Amazon MQ broker.

Topics

For more information about how Amazon MQ encrypts your data, as well as a list of supported protocols, see Data Protection.

Prefer brokers without public accessibility

Brokers created without public accessibility can't be accessed from outside of your VPC. This greatly reduces your broker's susceptibility to Distributed Denial of Service (DDoS) attacks from the public internet. For more information, see Accessing the broker web console without public accessibility in this guide and How to Help Prepare for DDoS Attacks by Reducing Your Attack Surface on the AWS Security Blog.

Always configure an authorization map

Because ActiveMQ has no authorization map configured by default, any authenticated user can perform any action on the broker. Thus, it is a best practice to restrict permissions *by group*. For more information, see `authorizationEntry`.

> **Important**
>
> If you specify an authorization map which doesn't include the `activemq-webconsole` group, you can't use the ActiveMQ Web Console because the group isn't authorized to send messages to, or receive messages from, the Amazon MQ broker.

Block unnecessary protocols with VPC security groups

To improve security, you should restrict the connections of unnecessary protocols and ports by properly configuring your Amazon VPC Security Group. For instance, to restrict access to most protocols while allowing access to OpenWire and the web console, you could allow access to only 61617 and 8162. This limits your exposure by blocking protocols you are not using, while allowing OpenWire and the web console to function normally.

Allow only the protocol ports that you are using.

- AMQP: 5671
- MQTT: 8883
- OpenWire: 61617
- STOMP: 61614
- WebSocket: 61619

For more information see:

- Configure Additional Broker Settings
- Security Groups for your VPC
- Default Security Group for Your VPC
- Working with Security Groups

Logging and monitoring Amazon MQ brokers

Monitoring is an important part of maintaining the reliability, availability, and performance of your AWS solutions. You should collect monitoring data from all of the parts of your AWS solution so that you can more easily debug a multi-point failure if one occurs. AWS provides several tools for monitoring your Amazon MQ resources and responding to potential incidents:

Topics

Accessing CloudWatch metrics for Amazon MQ

Amazon MQ and Amazon CloudWatch are integrated so you can use CloudWatch to view and analyze metrics for your ActiveMQ broker and the broker's destinations (queues and topics). You can view and analyze your Amazon MQ metrics from the CloudWatch console, the AWS CLI, or the CloudWatch CLI. CloudWatch metrics for Amazon MQ are automatically polled from the broker and then pushed to CloudWatch every minute.

For a full list of Amazon MQ metrics, see Monitoring Amazon MQ using CloudWatch.

For information about creating a CloudWatch alarm for a metrics, see Create or Edit a CloudWatch Alarm in the *Amazon CloudWatch User Guide*.

> **Note**
>
> There is no charge for the Amazon MQ metrics reported in CloudWatch. These metrics are provided as part of the Amazon MQ service.
> For ActiveMQ brokers, CloudWatch monitors only the first 1000 destinations.
> For RabbitMQ brokers, CloudWatch monitors only the first 500 destinations, ordered by number of consumers.

Topics

- AWS Management Console
- AWS Command Line Interface
- Amazon CloudWatch API

AWS Management Console

The following example shows you how to access CloudWatch metrics for Amazon MQ using the AWS Management Console.

> **Note**
>
> If you're already signed into the Amazon MQ console, on the broker **Details** page, choose **Actions**, **View CloudWatch metrics**.
>
> 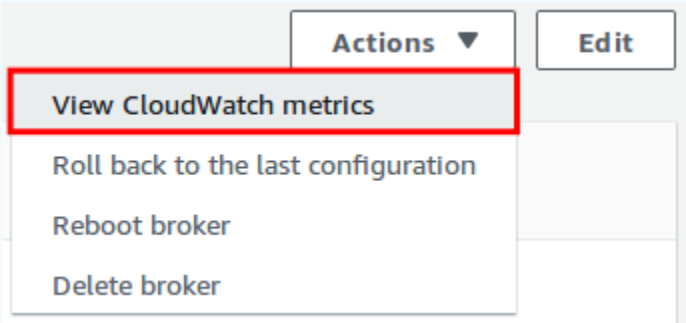
>

1. Sign in to the CloudWatch console.
2. On the navigation panel, choose **Metrics**.
3. Select the **AmazonMQ** metric namespace.

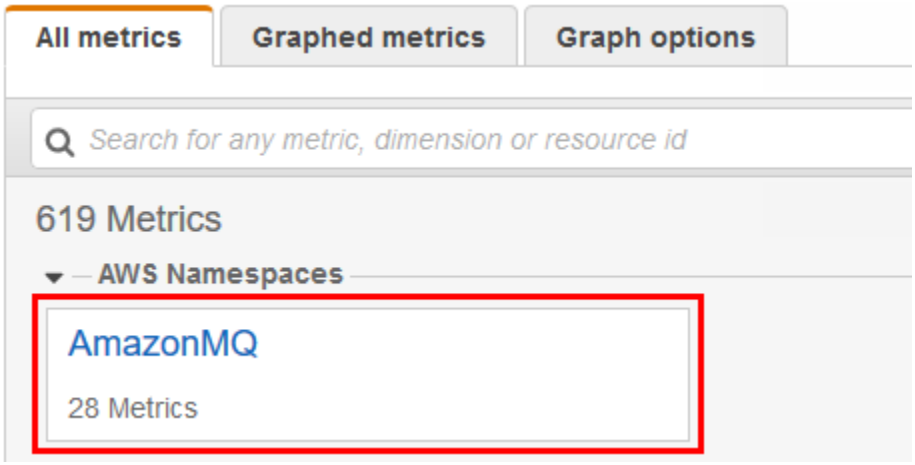

4. Select one of the following metric dimensions:

   - **Broker Metrics**
   - **Queue Metrics by Broker**
   - **Topic Metrics by Broker**

   In this example, **Broker Metrics** is selected.

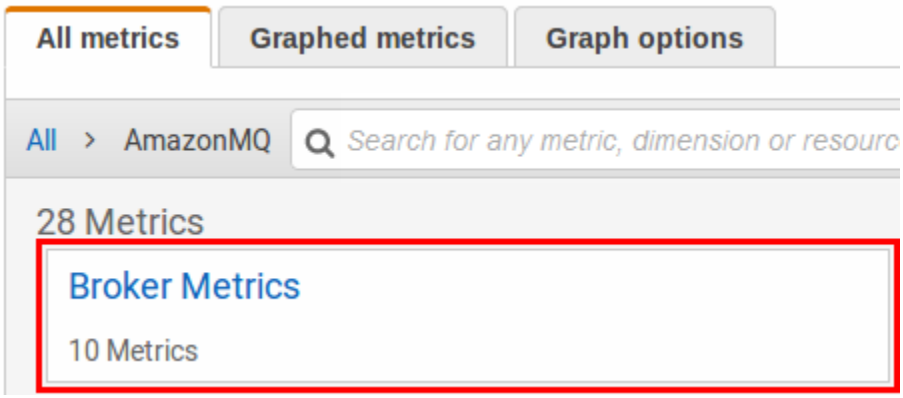

5. You can now examine your Amazon MQ metrics:

   - To sort the metrics, use the column heading.
   - To graph the metric, select the check box next to the metric.
   - To filter by metric, choose the metric name and then choose **Add to search**.

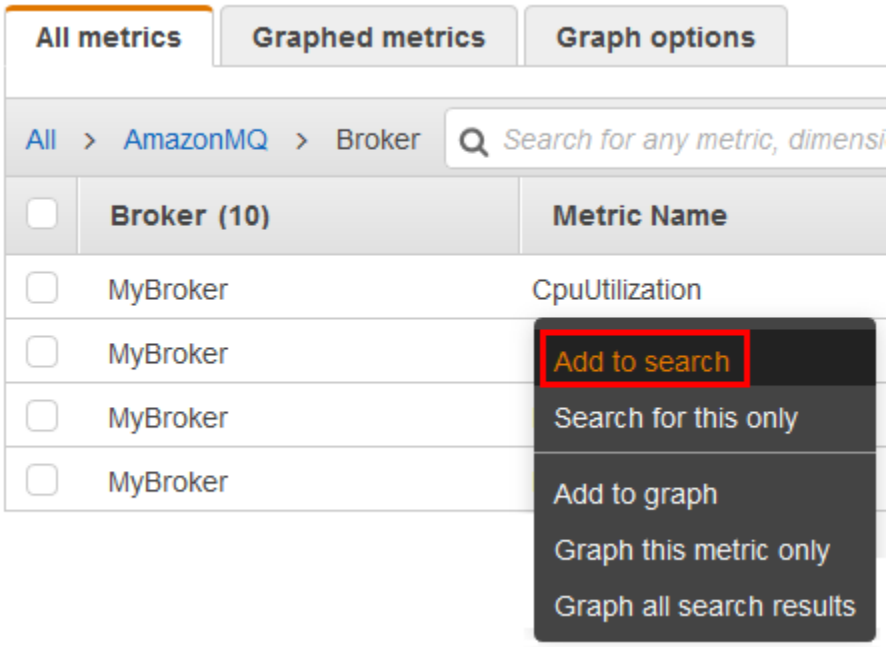

AWS Command Line Interface

To access Amazon MQ metrics using the AWS CLI, use the `get-metric-statistics` command.

For more information, see Get Statistics for a Metric in the *Amazon CloudWatch User Guide*.

Amazon CloudWatch API

To access Amazon MQ metrics using the CloudWatch API, use the `GetMetricStatistics` action.

For more information, see Get Statistics for a Metric in the *Amazon CloudWatch User Guide*.

Monitoring Amazon MQ brokers using Amazon CloudWatch

Amazon MQ and Amazon CloudWatch are integrated so you can use CloudWatch to view and analyze metrics for your ActiveMQ broker and the broker's destinations (queues and topics). You can view and analyze your Amazon MQ metrics from the CloudWatch console, the AWS CLI, or the CloudWatch CLI. CloudWatch metrics for Amazon MQ are automatically polled from the broker and then pushed to CloudWatch every minute.

For information, see Accessing CloudWatch metrics for Amazon MQ.

> **Note**
>
> The following statistics are valid for all of the metrics:
>
> - `Average`
> - `Minimum`
> - `Maximum`
> - `Sum`

The `AWS/AmazonMQ` namespace includes the following metrics.

Topics

- Logging and monitoring Amazon MQ for ActiveMQ brokers
- Logging and monitoring Amazon MQ for RabbitMQ brokers

Logging and monitoring Amazon MQ for ActiveMQ brokers

Amazon MQ for ActiveMQ metrics

Metric	Unit	Description
`AmqpMaximumConnections`	Count	The maximum number of clients you can connect to your broker using AMQP. For more information on connection quotas, see Quotas in Amazon MQ.
`BurstBalance`	Percent	The percentage of burst credits remaining on the Amazon EBS volume used to persist message data for throughput-optimized brokers. If this balance reaches zero, the IOPS

Metric	Unit	Description
		provided by the Amazon EBS volume will decrease until the Burst Balance refills. For more information on how Burst Balances work in Amazon EBS, see: I/O Credits and Burst Performance.

Metric	Unit	Description
`CpuCreditBalance`	Credits (vCPU-minutes)	**Important** This metric is available only for the `mq.t2.micro` broker instance type. CPU credit metrics are available only at five-minute intervals. The number of earned CPU credits that an instance has accrued since it was launched or started (including the number of launch credits). The credit balance is available for the broker instance to spend on bursts beyond the baseline CPU utilization. Credits are accrued in the credit balance after they're earned and removed from the credit balance after they're spent. The credit balance has a maximum limit. Once the limit is reached, any newly earned credits are discarded.
`CpuUtilization`	Percent	The percentage of allocated Amazon EC2 compute units that the broker currently uses.

Metric	Unit	Description
`CurrentConnectionsCount`	Count	The current number of active connections on the current broker.
`EstablishedConnectionsCount`	Count	The total number of connections, active and inactive, that have been established on the broker.
`HeapUsage`	Percent	The percentage of the ActiveMQ JVM memory limit that the broker currently uses.
`InactiveDurableTopicSubscribersCount`	Count	The number of inactive durable topic subscribers, up to a maximum of 2000.
`JobSchedulerStorePercentUsage`	Percent	The percentage of disk space used by the job scheduler store.
`JournalFilesForFastRecovery`	Count	The number of journal files that will be replayed after a clean shutdown.
`JournalFilesForFullRecovery`	Count	The number of journal files that will be replayed after an unclean shutdown.
`MqttMaximumConnections`	Count	The maximum number of clients you can connect to your broker using MQTT. For more information on connection quotas, see Quotas in Amazon MQ.

Metric	Unit	Description
`NetworkConnectorConnectionCount`	Count	The number of nodes connected to the broker in a network of brokers using NetworkConnector.
`NetworkIn`	Bytes	The volume of incoming traffic for the broker.
`NetworkOut`	Bytes	The volume of outgoing traffic for the broker.
`OpenTransactionCount`	Count	The total number of transactions in progress.
`OpenwireMaximumConnections`	Count	The maximum number of clients you can connect to your broker using OpenWire. For more information on connection quotas, see Quotas in Amazon MQ.
`StompMaximumConnections`	Count	The maximum number of clients you can connect to your broker using STOMP. For more information on connection quotas, see Quotas in Amazon MQ.
`StorePercentUsage`	Percent	The percent used by the storage limit. If this reaches 100, the broker will refuse messages.
`TempPercentUsage`	Percent	The percentage of available temporary storage used by non-persistent messages.

Metric	Unit	Description
`TotalConsumerCount`	Count	The number of message consumers subscribed to destinations on the current broker.
`TotalMessageCount`	Count	The number of messages stored on the broker.
`TotalProducerCount`	Count	The number of message producers active on destinati ons on the current broker.
`VolumeReadOps`	Count	The number of read operation s performed on the Amazon EBS volume.
`VolumeWriteOps`	Count	The number of write operations performed on the Amazon EBS volume.
`WsMaximumConnections`	Count	The maximum number of clients you can connect to your broker using WebSocket . For more information on connection quotas, see Quotas in Amazon MQ.

Dimensions for ActiveMQ broker metrics

Dimension	Description
`Broker`	The name of the broker

Dimension	Description
	Note A single-instance broker has the suffix -1. An active/standby broker for high availability has the suffixes -1 and -2 for its redundant pair.

ActiveMQ destination (queue and topic) metrics

Important

The following metrics include per-minute counts for the CloudWatch polling period.

- `EnqueueCount`
- `ExpiredCount`
- `DequeueCount`
- `DispatchCount`
- `InFlightCount`

For example, in a five-minute CloudWatch period, `EnqueueCount` has five count values, each for a one-minute portion of the period. The `Minimum` and `Maximum` statistics provide the lowest and highest per-minute value during the specified period.

Metric	Unit	Description
`ConsumerCount`	Count	The number of consumers subscribed to the destination.
`EnqueueCount`	Count	The number of messages sent to the destination, per minute.

Metric	Unit	Description
`EnqueueTime`	Time (milliseconds)	The end-to-end latency from when a message arrives at a broker until it is delivered to a consumer. **Note** `EnqueueTime` does not measure the end-to-end latency from when a message is sent by a producer until it reaches the broker, nor the latency from when a message is received by a broker until it is acknowled ged by the broker. Rather, `EnqueueTi me` is the number of milliseconds from the moment a message is received by the broker until it is successfu lly delivered to a consumer.
`ExpiredCount`	Count	The number of messages that couldn't be delivered because they expired, per minute.
`DispatchCount`	Count	The number of messages sent to consumers, per minute.

Metric	Unit	Description
`DequeueCount`	Count	The number of messages acknowledged by consumers, per minute.
`InFlightCount`	Count	The number of messages sent to consumers that have not been acknowledged.
`ReceiveCount`	Count	The number of messages that have been received from the remote broker for a duplex network connector.
`MemoryUsage`	Percent	The percentage of the memory limit that the destination currently uses.
`ProducerCount`	Count	The number of producers for the destination.
`QueueSize`	Count	The number of messages in the queue. **Important** This metric applies only to queues.
`TotalEnqueueCount`	Count	The total number of messages that have been sent to the broker.
`TotalDequeueCount`	Count	The total number of messages that have been consumed by clients.

> **Note**
>
> `TotalEnqueueCount` and `TotalDequeueCount` metrics include messages for advisory topics. For more information about advisory topic messages, see the ActiveMQ documentation.

Dimensions for ActiveMQ destination (queue and topic) metrics

Dimension	Description
`Broker`	The name of the broker. **Note** A single-instance broker has the suffix `-1`. An active/standby broker for high availability has the suffixes `-1` and `-2` for its redundant pair.
`Topic` or `Queue`	The name of the topic or queue.
`NetworkConnector`	The name of the network connector.

Logging and monitoring Amazon MQ for RabbitMQ brokers

RabbitMQ broker metrics

Metric	Unit	Description
`ExchangeCount`	Count	The total number of exchanges configured on the broker.
`QueueCount`	Count	The total number of queues configured on the broker.

Metric	Unit	Description
`ConnectionCount`	Count	The total number of connections established on the broker.
`ChannelCount`	Count	The total number of channels established on the broker.
`ConsumerCount`	Count	The total number of consumers connected to the broker.
`MessageCount`	Count	The total number of messages in the queues. **Note** The number produced is the total sum of ready and unacknowl edged messages on the broker.
`MessageReadyCount`	Count	The total number of ready messages in the queues.
`MessageUnacknowled gedCount`	Count	The total number of unacknowledged messages in the queues.
`PublishRate`	Count	The rate at which messages are published to the broker. The number produced represents the number of messages per second at the time of sampling.

Metric	Unit	Description
`ConfirmRate`	Count	The rate at which the RabbitMQ server is confirming published messages. You can compare this metric with `PublishRate` to better understand how your broker is performing. The number produced represents the number of messages per second at the time of sampling.
`AckRate`	Count	The rate at which messages are being acknowledged by consumers. The number produced represents the number of messages per second at the time of sampling.
`SystemCpuUtilization`	Percent	The percentage of allocated Amazon EC2 compute units that the broker currently uses. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.

Metric	Unit	Description
`RabbitMQMemLimit`	Bytes	The RAM limit for a RabbitMQ broker. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.
`RabbitMQMemUsed`	Bytes	The volume of RAM used by a RabbitMQ broker. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values.
`RabbitMQDiskFreeLimit`	Bytes	The disk limit for a RabbitMQ broker. For cluster deployments, this value represents the aggregate of all three RabbitMQ nodes' corresponding metric values. This metric is different per instance size. For more information about Amazon MQ instance types, see the section called “Amazon MQ for RabbitMQ instance types”.

Metric	Unit	Description
`RabbitMQDiskFree`	Bytes	The total volume of free disk space available in a RabbitMQ broker. When disk usage goes above its limit, the cluster will block all producer connectio ns. For cluster deploymen ts, this value represents the aggregate of all three RabbitMQ nodes' correspon ding metric values.
`RabbitMQFdUsed`	Count	Number of file descriptors used. For cluster deploymen ts, this value represents the aggregate of all three RabbitMQ nodes' correspon ding metric values.
`RabbitMQIOReadAver ageTime`	Count	The average time (in milliseco nds) for RabbitMQ to perform one read operation. The value is proportional to the message size.
`RabbitMQIOWriteAve rageTime`	Count	The average time (in milliseco nds) for RabbitMQ to perform one write operation. The value is proportional to the message size.

Dimensions for RabbitMQ broker metrics

Dimension	Description
`Broker`	The name of the broker.

RabbitMQ node metrics

Metric	Unit	Description
`SystemCpuUtilization`	Percent	The percentage of allocated Amazon EC2 compute units that the broker currently uses.
`RabbitMQMemLimit`	Bytes	The RAM limit for a RabbitMQ node.
`RabbitMQMemUsed`	Bytes	The volume of RAM used by a RabbitMQ node. When memory use goes above the limit, the cluster will block all producer connections.
`RabbitMQDiskFreeLimit`	Bytes	The disk limit for a RabbitMQ node. This metric is different per instance size. For more information about Amazon MQ instance types, see the section called "Amazon MQ for RabbitMQ instance types".
`RabbitMQDiskFree`	Bytes	The total volume of free disk space available in a RabbitMQ node. When disk usage goes above its limit, the cluster will block all producer connections.

Metric	Unit	Description
`RabbitMQFdUsed`	Count	Number of file descriptors used.

Dimensions for RabbitMQ node metrics

Dimension	Description
`Node`	The name of the node. **Note** A node name consists of two parts: a prefix (usuallly `rabbit`) and a hostname. For example, `rabbit@ip-10-0-0-230.us-west-2.compute.internal` is a node name with the prefix `rabbit` and the hostname `ip-10-0-0-230.us-west-2.compute.internal`.
`Broker`	The name of the broker.

RabbitMQ queue metrics

Metric	Unit	Description
`ConsumerCount`	Count	The number of consumers subscribed to the queue.
`MessageReadyCount`	Count	The number of messages that are currently available to be delivered.

Metric	Unit	Description
`MessageUnacknowledgedCount`	Count	The number of messages for which the server is awaiting acknowledgement.
`MessageCount`	Count	The total number of `MessageReadyCount` and `MessageUnacknowledgedCount` (also known as queue depth).

Dimensions for RabbitMQ queue metrics

> **Note**
>
> Amazon MQ for RabbitMQ will not publish metrics for virtual hosts and queues with names containing blank spaces, tabs or other non-ASCII characters.
> For more information about dimension names, see Dimension in the *Amazon CloudWatch API Reference*.

Dimension	Description
`Queue`	The name of the queue.
`VirtualHost`	The name of the virtual host.
`Broker`	The name of the broker.

Logging Amazon MQ API calls using AWS CloudTrail

Amazon MQ is integrated with AWS CloudTrail, a service that provides a record of the Amazon MQ calls that a user, role, or AWS service makes. CloudTrail captures API calls related to Amazon MQ brokers and configurations as events, including calls from the Amazon MQ console and code calls from Amazon MQ APIs. For more information about CloudTrail, see the *AWS CloudTrail User Guide*.

> **Note**
>
> CloudTrail doesn't log API calls related to ActiveMQ operations (for example, sending and receiving messages) or to the ActiveMQ Web Console. To log information related to ActiveMQ operations, you can configure Amazon MQ to publish general and audit logs to Amazon CloudWatch Logs.

Using the information that CloudTrail collects, you can identify a specific request to an Amazon MQ API, the IP address of the requester, the requester's identity, the date and time of the request, and so on. If you configure a *trail*, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket. If you don't configure a trail, you can view the most recent events in the event history in the CloudTrail console. For more information, see Overview for Creating a Trail in the *AWS CloudTrail User Guide*.

Amazon MQ Information in CloudTrail

When you create your AWS account, CloudTrail is enabled. When a supported Amazon MQ event activity occurs, it is recorded in a CloudTrail event with other AWS service events in the event history. You can view, search, and download recent events for your AWS account. For more information, see Viewing Events with CloudTrail Event History in the *AWS CloudTrail User Guide*.

A trail allows CloudTrail to deliver log files to an Amazon S3 bucket. You can create a trail to keep an ongoing record of events in your AWS account. By default, when you create a trail using the AWS Management Console, the trail applies to all AWS Regions. The trail logs events from all AWS Regions and delivers log files to the specified Amazon S3 bucket. You can also configure other AWS services to further analyze and act on the event data collected in CloudTrail logs. For more information, see the following topics in the *AWS CloudTrail User Guide*:

- CloudTrail Supported Services and Integrations
- Configuring Amazon SNS Notifications for CloudTrail
- Receiving CloudTrail Log Files from Multiple Regions
- Receiving CloudTrail Log Files from Multiple Accounts

Amazon MQ supports logging both the request parameters and the responses for the following APIs as events in CloudTrail log files:

- `CreateConfiguration`
- `DeleteBroker`
- `DeleteUser`
- `RebootBroker`
- `UpdateBroker`

> **Note**
>
> RebootBroker log files are logged when you reboot the broker. During the maintenance window, the service automatically reboots, and RebootBroker log files are not logged.

> **Important**
>
> For the `GET` methods of the following APIs, the request parameters are logged, but the responses are redacted:
>
> - `DescribeBroker`
> - `DescribeConfiguration`
> - `DescribeConfigurationRevision`
> - `DescribeUser`
> - `ListBrokers`
> - `ListConfigurationRevisions`
> - `ListConfigurations`
> - `ListUsers`
>
> For the following APIs, the `data` and `password` request parameters are hidden by asterisks (`***`):
>
> - `CreateBroker` (POST)
> - `CreateUser` (POST)
> - `UpdateConfiguration` (PUT)
> - `UpdateUser` (PUT)

Every event or log entry contains information about the requester. This information helps you determine the following:

- Was the request made with root or user credentials?
- Was the request made with temporary security credentials for a role or a federated user?
- Was the request made by another AWS service?

For more information, see CloudTrail userIdentity Element in the *AWS CloudTrail User Guide*.

Example Amazon MQ Log File Entry

A *trail* is a configuration that allows the delivery of events as log files to the specified Amazon S3 bucket. CloudTrail log files contain one or more log entries.

An *event* represents a single request from any source and includes information about the request to an Amazon MQ API, the IP address of the requester, the requester's identity, the date and time of the request, and so on.

The following example shows a CloudTrail log entry for a `CreateBroker` API call.

> **Note**
>
> Because CloudTrail log files aren't an ordered stack trace of public APIs, they don't list information in any specific order.

```
{
    "eventVersion": "1.06",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "AKIAIOSFODNN7EXAMPLE",
        "arn": "arn:aws:iam::111122223333:user/AmazonMqConsole",
        "accountId": "111122223333",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "userName": "AmazonMqConsole"
    },
    "eventTime": "2018-06-28T22:23:46Z",
    "eventSource": "amazonmq.amazonaws.com",
    "eventName": "CreateBroker",
    "awsRegion": "us-west-2",
```

```
    "sourceIPAddress": "203.0.113.0",
    "userAgent": "PostmanRuntime/7.1.5",
    "requestParameters": {
        "engineVersion": "5.15.9",
        "deploymentMode": "ACTIVE_STANDBY_MULTI_AZ",
        "maintenanceWindowStartTime": {
            "dayOfWeek": "THURSDAY",
            "timeOfDay": "22:45",
            "timeZone": "America/Los_Angeles"
        },
        "engineType": "ActiveMQ",
        "hostInstanceType": "mq.m5.large",
        "users": [
            {
                "username": "MyUsername123",
                "password": "***",
                "consoleAccess": true,
                "groups": [
                    "admins",
                    "support"
                ]
            },
            {
                "username": "MyUsername456",
                "password": "***",
                "groups": [
                    "admins"
                ]
            }
        ],
        "creatorRequestId": "1",
        "publiclyAccessible": true,
        "securityGroups": [
            "sg-a1b234cd"
        ],
        "brokerName": "MyBroker",
        "autoMinorVersionUpgrade": false,
        "subnetIds": [
            "subnet-12a3b45c",
            "subnet-67d8e90f"
        ]
    },
    "responseElements": {
        "brokerId": "b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9",
```

```
        "brokerArn": "arn:aws:mq:us-
east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9"
    },
    "requestID": "a1b2c345-6d78-90e1-f2g3-4hi56jk7l890",
    "eventID": "a12bcd3e-fg45-67h8-ij90-12k34d5l16mn",
    "readOnly": false,
    "eventType": "AwsApiCall",
    "recipientAccountId": "111122223333"
}
```

Configuring Amazon MQ to publish logs to Amazon CloudWatch Logs

Amazon MQ is integrated with Amazon CloudWatch Logs, a service that monitors, stores, and accesses your log files from a variety of sources. For example, you can configure CloudWatch alarms to receive notifications of broker reboots or troubleshoot ActiveMQ broker configuration errors. For more information about CloudWatch Logs, see the *Amazon CloudWatch Logs User Guide*

Topics

- Configuring Amazon MQ for ActiveMQ logs
- Configuring Amazon MQ for RabbitMQ logs

Configuring Amazon MQ for ActiveMQ logs

To allow Amazon MQ to publish logs to CloudWatch Logs, you must add a permission to your Amazon MQ user and also configure a resource-based policy for Amazon MQ before you create or restart the broker.

The following describes the steps to configure CloudWatch logs for your ActiveMQ brokers.

Topics

- Understanding the structure of logging in CloudWatch Logs
- Add the CreateLogGroup permission to your Amazon MQ user
- Configure a resource-based policy for Amazon MQ
- Cross-service confused deputy prevention
- Troubleshooting CloudWatch Logs Configuration

Understanding the structure of logging in CloudWatch Logs

You can enable *general* and *audit* logging when you configure advanced broker settings when you create a broker, or when you edit a broker.

General logging enables the default `INFO` logging level (`DEBUG` logging isn't supported) and publishes `activemq.log` to a log group in your CloudWatch account. The log group has a format similar to the following:

```
/aws/amazonmq/broker/b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9/general
```

Audit logging enables logging of management actions taken using JMX or using the ActiveMQ Web Console and publishes `audit.log` to a log group in your CloudWatch account. The log group has a format similar to the following:

```
/aws/amazonmq/broker/b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9/audit
```

Depending on whether you have a single-instance broker or an active/standby broker, Amazon MQ creates either one or two log streams within each log group. The log streams have a format similar to the following.

```
activemq-b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.log
activemq-b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-2.log
```

The `-1` and `-2` suffixes denote individual broker instances. For more information, see Working with Log Groups and Log Streams in the *Amazon CloudWatch Logs User Guide*.

Add the `CreateLogGroup` permission to your Amazon MQ user

To allow Amazon MQ to create a CloudWatch Logs log group, you must ensure that the user who creates or reboots the broker has the `logs:CreateLogGroup` permission.

> **Important**
>
> If you don't add the `CreateLogGroup` permission to your Amazon MQ user before the user creates or reboots the broker, Amazon MQ doesn't create the log group.

The following example IAM-based policy grants permission for `logs:CreateLogGroup` for users to whom this policy is attached.

```
{
   "Version": "2012-10-17",
   "Statement": [
      {
         "Effect": "Allow",
         "Action": "logs:CreateLogGroup",
         "Resource": "arn:aws:logs:*:*:log-group:/aws/amazonmq/*"
      }
   ]
}
```

> **Note**
>
> Here, the term user refers to *Users* and not *Amazon MQ users*, which are created when a new broker is configured. For more information regarding setting up users and configuring IAM policies, please refer to the Identity Management Overview section of the IAM User Guide.

For more information, see `CreateLogGroup` in the *Amazon CloudWatch Logs API Reference*.

Configure a resource-based policy for Amazon MQ

> **Important**
>
> If you don't configure a resource-based policy for Amazon MQ, the broker can't publish the logs to CloudWatch Logs.

To allow Amazon MQ to publish logs to your CloudWatch Logs log group, configure a resource-based policy to give Amazon MQ access to the following CloudWatch Logs API actions:

- `CreateLogStream` – Creates a CloudWatch Logs log stream for the specified log group.
- `PutLogEvents` – Delivers events to the specified CloudWatch Logs log stream.

The following resource-based policy grants permission for `logs:CreateLogStream` and `logs:PutLogEvents` to AWS.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": { "Service": "mq.amazonaws.com" },
            "Action": [ "logs:CreateLogStream", "logs:PutLogEvents" ],
            "Resource": "arn:aws:logs:*:*:log-group:/aws/amazonmq/*"
        }
    ]
}
```

This resource-based policy *must* be configured by using the AWS CLI as shown by the following command. In the example, replace *`us-east-1`* with your own information.

```
aws --region us-east-1 logs put-resource-policy --policy-name AmazonMQ-logs \
--policy-document "{\"Version\": \"2012-10-17\", \"Statement\":[{ \"Effect\": \"Allow
\", \"Principal\": { \"Service\": \"mq.amazonaws.com\" },
\"Action\": [\"logs:CreateLogStream\", \"logs:PutLogEvents\"], \"Resource\":
 \"arn:aws:logs:*:*:log-group:\/aws\/amazonmq\/*\" }]}"
```

> **Note**
>
> Because this example uses the `/aws/amazonmq/` prefix, you need to configure the resource-based policy only once per AWS account, per region.

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that help you protect

your data for all services with service principals that have been given access to resources in your account.

We recommend using the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in your Amazon MQ resource-based policy to limit CloudWatch Logs access to one or more specified brokers.

> **Note**
>
> If you use both global condition context keys, the `aws:SourceAccount` value and the account in the `aws:SourceArn` value must use the same account ID when used in the same policy statement.

The following example demonstrates a resource-based policy that limits CloudWatch Logs access to a single Amazon MQ broker.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "mq.amazonaws.com"
      },
      "Action": [
        "logs:CreateLogStream",
        "logs:PutLogEvents"
      ],
      "Resource": "arn:aws:logs:*:*:log-group:/aws/amazonmq/*",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "123456789012",
          "aws:SourceArn": "arn:aws:mq:us-
east-2:123456789012:broker:MyBroker:b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9"
        }
      }
    }
  ]
}
```

You can also configure your resource-based policy to limit CloudWatch Logs access to all brokers in an account, as shown in the following.

```
{
    "Version": "2012-10-17",
    "Statement": [
      {
        "Effect": "Allow",
        "Principal": {
          "Service": [
            "mq.amazonaws.com"
          ]
        },
        "Action": [
          "logs:CreateLogStream",
          "logs:PutLogEvents"
        ],
        "Resource": "arn:aws:logs:*:*:log-group:/aws/amazonmq/*",
        "Condition": {
          "ArnLike": {
            "aws:SourceArn": "arn:aws:mq:*:123456789012:broker:*"
          },
          "StringEquals": {
            "aws:SourceAccount": "123456789012"
          }
        }
      }
    ]
  }
```

For more information about the confused deputy security issue, see The confused deputy problem in the *User Guide*.

Troubleshooting CloudWatch Logs Configuration

In some cases, CloudWatch Logs might not always behave as expected. This section gives an overview of common issues and shows how to resolve them.

Log Groups Don't Appear in CloudWatch

Add the `CreateLogGroup` permission to your Amazon MQ user and reboot the broker. This allows Amazon MQ to create the log group.

Log Streams Don't Appear in CloudWatch Log Groups

Configure a resource-based policy for Amazon MQ. This allows your broker to publish its logs.

Configuring Amazon MQ for RabbitMQ logs

When you enable CloudWatch logging for your RabbitMQ brokers, Amazon MQ uses a service-linked role to publish general logs to CloudWatch. If no Amazon MQ service-linked role exists when you first create a broker, Amazon MQ will automatically create one. All subsequent RabbitMQ brokers will use the same service-linked role to publish logs to CloudWatch.

For more information about service-linked roles, see Using service-linked roles in the *AWS Identity and Access Management User Guide*. For more information about how Amazon MQ uses service-linked roles, see the section called “Using service-linked roles”.

Quotas in Amazon MQ

This topic lists quotas within Amazon MQ. Many of the following quotas can be changed for specific AWS accounts. To request an increase for a limit, see AWS Service Quotas in the *Amazon Web Services General Reference*. Updated limits will not be visible even after the limit increase has been applied. For more information on viewing current connection limits in Amazon CloudWatch, see Monitoring Amazon MQ brokers using Amazon CloudWatch.

Topics

Brokers

The following table lists quotas related to Amazon MQ brokers.

Limit	Description
Broker name	• Must be unique in your AWS account. • Must be 1-50 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~).
Number of brokers, per region	50

Limit	Description
Wire-level connections per protocol for smaller broker	**Important** Does not apply to RabbitMQ brokers. 300 for `mq.*.micro` instance type brokers.
Wire-level connections per protocol for larger broker	**Important** Does not apply to RabbitMQ brokers. 2,000 for `mq.*.*large` instance type brokers.
Security groups per broker	5
ActiveMQ destinations (queues, and topics) monitored in CloudWatch	CloudWatch monitors only the first 1000 destinations.
RabbitMQ destinations (queues) monitored in CloudWatch	CloudWatch monitors only the first 500 destinations, ordered by number of consumers .
Tags per broker	50

Configurations

The following table lists quotas related to Amazon MQ configurations.

> **Important**
> Does not apply to RabbitMQ brokers.

Limit	Description
Configuration name	• Must be 1-150 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~).
Revisions per configuration	300

Users

The following table lists quotas related to Amazon MQ ActiveMQ broker users.

> **Important**
>
> Does not apply to RabbitMQ brokers.

Limit	Description
Username	• Must be 1-100 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Can contain only alphanumeric characters, dashes, periods, underscores, and tildes (- . _ ~). • Must not contain commas (,).
Password	•

Limit	Description
	Must be 12-250 characters long. • Must contain only characters specified in the ASCII Printable Character Set. • Must contain at least 4 unique characters. • Must not contain commas (`,`).
Users per broker (simple auth)	250
Groups per user (simple auth)	20

Data Storage

The following table lists quotas related to Amazon MQ data storage.

Limit	Description
Storage capacity per smaller broker	20 GB for `mq.*.micro` instance type brokers. For more information regarding A mazon MQ instance types, see Broker instance types.
Storage capacity per larger broker	200 GB for `mq.*.*large` instance type brokers. For more information regarding A mazon MQ instance types, see Broker instance types.
Job scheduler usage limit per broker backed by Amazon EBS	**Important** Does not apply to RabbitMQ brokers.

Limit	Description
	50 GB. For more information about job scheduler usage, see `JobSchedulerUsage` in the *Apache ActiveMQ API Documentation*.
Temporary storage capacity per smaller broker.	**Important** Does not apply to RabbitMQ brokers. 5 GB for `mq.*.micro` instance type brokers.
Temporary storage capacity per larger broker.	**Important** Does not apply to RabbitMQ brokers. 50 GB for `mq.*.*large` instance type brokers.

API Throttling

The following throttling quotas are aggregated per AWS account, *across all Amazon MQ APIs* to maintain service bandwidth. For more information about Amazon MQ APIs, see the *Amazon MQ REST API Reference*.

> **Important**
>
> These quotas don't apply to Amazon MQ for ActiveMQ or Amazon MQ for RabbitMQ broker messaging APIs. For example, Amazon MQ doesn't throttle the sending or receiving of messages.

API burst limit	API rate limit
100	15

Troubleshooting Amazon MQ

This section describes common issues you might encounter when using Amazon MQ brokers, and the steps you can take to resolve them.

Contents

Troubleshooting: General

Use the information in this section to help you diagnose common issues you might encounter when working with Amazon MQ brokers, such as issues connecting to your broker, and broker reboots.

Contents

I can't connect to my broker web console or endpoints.

If you're experiencing issues connecting to your broker using the web console or wire-level endpoints, we recommend the following steps.

1. Check whether you're attempting to connect to your broker from behind a firewall. You might need to configure the firewall to allow access to your broker.
2. Check whether you're trying to connect to your broker using a FIPS endpoint. Amazon MQ only supports FIPS endpoints when using API operations, and not for wire-level connections to the broker instance itself.
3. Check if the **Public Accessibility** option for your broker is set to **Yes**. If this is set to **No**, check your subnet's network Access Control List (ACL) rules. If you've created custom network ACLs, you might need to change the network ACL rules to provide access to your broker. For more

information about Amazon VPC networking, see Enabling internet access in the *Amazon VPC User Guide*

4. Check your broker's Security Group rules. Make sure that you are allowing connections to the following ports:

> **Note**
>
> The following ports are grouped according to engine types because Amazon MQ for ActiveMQ and Amazon MQ for RabbitMQ use different ports for connections.

Amazon MQ for ActiveMQ

- Web console – Port 8162
- OpenWire – Port 61617
- AMQP – Port 5671
- STOMP – Port 61614
- MQTT – Port 8883
- WSS – Port 61619

Amazon MQ for RabbitMQ

- Web console and management API – Port 443 and 15671
- AMQP – Port 5671

5. Run the following network connectivity tests for your broker engine type.

> **Note**
>
> For brokers without public accessibility, run the tests from an Amazon EC2 instance within the same Amazon VPC as your Amazon MQ broker and evaluate the responses.

Amazon MQ for ActiveMQ

To test your Amazon MQ for ActiveMQ broker's network connectivity

1. Open a new terminal or command line window.

2. Run the following `nslookup` command to query your broker DNS record. For active/standby deployments, test both the active and standby endpoints. The active/standby endpoints are identified with a suffix, `-1` or `-2` added to the unique broker ID. Replace the endpoint with your information.

```
$ nslookup b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-west-2.amazonaws.com
```

If the query succeeds, you will see an output similar to the following.

```
Non-authoritative answer:
Server:  dns-resolver-corp-sfo-1.sfo.corp.amazon.com
Address:  172.10.123.456

Name:    ec2-12-345-123-45.us-west-2.compute.amazonaws.com
Address:  12.345.123.45
Aliases:  b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-west-2.amazonaws.com
```

The resolved IP address should match the IP addresses provided in the Amazon MQ console. This indicates that the domain name is resolving correctly on the DNS server, and you can move on to the next step.

3. Run the following `telnet` command to test the network path for your broker. Replace the endpoint with your information. Replace *`port`* with port number 8162 for the web console, or other wire-level ports to test additional protocols as needed.

> **Note**
>
> For active/standby deployments, you will receive a `Connect failed` error message if you run `telnet` with the standby endpoint. This is expected, as the standby instance itself is running, but the ActiveMQ process is not running and does not have access to the broker's Amazon EFS storage volume. Run the command for both `-1` and `-2` endpoints to ensure you test both the active and the standby instances.

```
$ telnet b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
west-2.amazonaws.com port
```

For the active instance, you will see an output similar to the following.

```
Connected to b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
west-2.amazonaws.com.
Escape character is '^]'.
```

4. Do one of the following.
   - If the `telnet` command succeeds, check the `EstablishedConnectionsCount` metric and confirm that the broker has not reached the maximum Wire-level connection limit. You can also confirm if the limit has been reached by reviewing the broker `General` logs. If this metric is greater than zero, then there is at least one client currently connected to the broker. If the metric shows zero connections, then perform the `telnet` path test again and wait at least one minute before disconnecting, as broker metrics are published every minute.
   - If the `telnet` command fails, check the status of your broker's elastic network interface, and confirm that the status is `in-use`. Create an Amazon VPC flow log for each instance's network interface, and review the generated flow logs. Look for the broker's IP addresses when you ran the `telnet` command, and confirm the connection packets are ACCEPTED, including a return packet. For more information, and to see a flow log example, see Flow log record examples in the *Amazon VPC Developer Guide*.
5. Run the following `curl` command to check connectivity to the ActiveMQ admin web console.

```
$ curl https://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
west-2.amazonaws.com:8162/index.html
```

If the command succeeds, the output should be an HTML document similar to the following.

```
<!DOCTYPE html PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN" "http://
www.w3.org/TR/html4/loose.dtd">
<html>
    <head>
        <meta http-equiv="Content-Type" content="text/html;
 charset=iso-8859-1" />
        <title>Apache ActiveMQ</title>
        ...
```

Amazon MQ for RabbitMQ

To test your Amazon MQ for RabbitMQ broker's network connectivity

1. Open a new terminal or command line window.
2. Run the following `nslookup` command to query your broker DNS record. Replace the endpoint with your information.

```
$ nslookup b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-west-2.amazonaws.com
```

If the query succeeds, you will see an output similar to the following.

```
Non-authoritative answer:
Server:  dns-resolver-corp-sfo-1.sfo.corp.amazon.com
Address:  172.10.123.456

Name:    rabbit-broker-1c23e456ca78-b9000123b4ebbab5.elb.us-
west-2.amazonaws.com
Addresses:  52.12.345.678
          52.23.234.56
          41.234.567.890
          54.123.45.678
Aliases:  b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-west-2.amazonaws.com
```

3. Run the following `telnet` command to test the network path for your broker. Replace the endpoint with your information. You can replace *`port`* with port 443 for the web console, and 5671 to test the wire-level AMQP connection.

```
$ telnet b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
west-2.amazonaws.com port
```

If the command succeeds, you'll see an output similar to the following.

```
Connected to b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
west-2.amazonaws.com.
Escape character is '^]'.
```

> **Note**
>
> The telnet connection will close automatically after a few seconds.

4. Do one of the following.
   - If the `telnet` command succeeds, check the `ConnectionCount` metric and confirm that the broker has not reached the value set in the `max-connections` default policy. You can also confirm if the limit has been reached by reviewing the broker `Connection.log` log group. If this metric is greater than zero, there is at least one client currently connected to the broker. If the metric shows zero connections, then perform the `telnet` path test again. You may need to repeat this process if the connection closes before your broker has published new connection metrics to CloudWatch. Metrics are published every minute.
   - For brokers without public accessibility, if the `telnet` command fails, check the status of your broker's elastic network interfaces, and confirm that the status is `in-use`. Create an Amazon VPC flow log for each network interface, and review the generated flow logs. Look for the broker's private IP addresses when you the `telnet` command was invoked, and confirm the connection packets are `ACCEPTED`, including a return packet. For more information, and to see a flow log example, see Flow log record examples in the *Amazon VPC Developer Guide*.

     > **Note**
     >
     > This step does not apply to Amazon MQ for RabbitMQ brokers with public accessibility.

5. Run the following `curl` command to check connectivity to the RabbitMQ admin web console.

```
$ curl https://b-1234a5b6-78cd-901e-2fgh-3i45j6k178l9-1.mq.us-
west-2.amazonaws.com:443/index.html
```

If the command succeeds, the output should be an HTML document similar to the following.

```
<!DOCTYPE html>
```

```
<html>
    <head>
        <meta http-equiv="X-UA-Compatible" content="IE=edge" />
        <meta http-equiv="Content-Type" content="text/html; charset=utf-8" />
        <title>RabbitMQ Management</title>
        ...
```

My broker is running, and I can verify connectivity using `telnet`, but my clients are unable to connect and are returning SSL exceptions.

Your broker endpoint certificate may have been updated during the broker maintenance window. Amazon MQ broker certificates are rotated periodically to ensure continued availability and security of brokers.

We recommend using the Amazon root certificate authority (CA) in Amazon Trust Services to authenticate against in your clients' trust store. All Amazon MQ broker certificates are signed with this root CA. By using an Amazon root CA, you will no longer need to download the new Amazon MQ broker certificate every time there is a certificate update on the broker.

I created a broker but broker creation failed.

If your broker is in a `CREATION_FAILED` status, do the following.

- Check your IAM permissions. To create a broker must either use the AWS managed IAM policy `AmazonMQFullAccess` or have the correct set of Amazon EC2 permissions in your custom IAM policy. To learn more about the required Amazon EC2 permissions you need, see IAM permissions required to create an Amazon MQ broker.
- Check if the subnet you are choosing for your broker is in a shared Amazon Virtual Private Cloud (VPC). To create an Amazon MQ broker in a shared Amazon VPC, you must create it in the account that owns the Amazon VPC.

My broker restarted and I'm not sure why.

If your broker has restarted automatically, it may be due to one of the following reasons.

- Your broker may have restarted because of a scheduled weekly maintenance window. Periodically, Amazon MQ performs maintenance to the hardware, operating system, or the

engine software of a message broker. The duration of the maintenance varies, but can last up to two hours, depending on the operations that are scheduled for your message broker. Brokers might restart at any point during the two hour maintenance window. For more information about broker maintenance windows, see the section called "Maintaining a broker".

- Your broker instance type might not be suitable to your application workload. For example, running a production workload on a `mq.t2.micro` might result in the broker running out of resources. High CPU utilization, or high broker memory usage can cause a broker to unexpectedly restart. To see how much CPU and memory is being utilized by your broker, use the following CloudWatch metrics for your engine type.
  - **Amazon MQ for ActiveMQ** – Check `CpuUtilization` for the percentage of allocated Amazon EC2 compute units that the broker currently uses. Check `HeapUsage`for the percentage of the ActiveMQ JVM memory limit that the broker currently uses.
  - **Amazon MQ for RabbitMQ** – Check `SystemCpuUtilization` for the percentage of allocated Amazon EC2 compute units that the broker currently uses. Check `RabbitMQMemUsed` for the volume of RAM used in Bytes, and divide by `RabbitMQMemLimit` for the percentage of memory used by the RabbitMQ node.

  For more information about broker instance types and how to choose the right instance type for your workload, see Broker instance types.

Troubleshooting: Amazon MQ for ActiveMQ

Use the information in this section to help you diagnose and resolve common issues you might encounter when working with Amazon MQ for ActiveMQ brokers.

Contents

I can't see general or audit logs for my broker in CloudWatch Logs even though I've activated logging.

If you're unable to view logs for your broker in CloudWatch Logs, do the following.

1. Check if the user who creates or reboots the broker has the `logs:CreateLogGroup` permission. If you don't add the `CreateLogGroup` permission to a user before the user creates or reboots the broker, Amazon MQ will not create the log group.
2. Check if you have configured a resource-based policy to allow Amazon MQ to publish logs to CloudWatch Logs. To allow Amazon MQ to publish logs to your CloudWatch Logs log group, configure a resource-based policy to give Amazon MQ access to the following CloudWatch Logs API actions:
   - `CreateLogStream` – Creates a CloudWatch Logs log stream for the specified log group.
   - `PutLogEvents` – Delivers events to the specified CloudWatch Logs log stream.

For more information about configuring Amazon MQ for ActiveMQ to publish logs to CloudWatch Logs, see Configuring logging.

After broker restart or maintenance window, I can't connect to my broker even though the status is `RUNNING`. Why?

You might be encountering connection issues after a broker restart you initiated, after a scheduled maintenance window is completed, or in a failure event, where the standby instance is activated. In either case, connection issues following a broker restart are most likely caused by unusually large numbers of messages persisted in your broker's Amazon EFS or Amazon EBS storage volume. During a restart, Amazon MQ moves persisted messages from storage to broker memory. To confirm this diagnosis, you can monitor the following metrics on CloudWatch for your Amazon MQ for ActiveMQ broker:

- **`StoragePercentUsage`** — Large percentages at or close to 100 percent can cause the broker to refuse connections.
- **`JournalFilesForFullRecovery`** — Indicates the number of of journal files that will be replayed following an unclean shutdown and restart. An increasing, or consistently higher than one, value indicates unresolved transactions that can cause connection issues after restart.

- **`OpenTransactionCount`** — A number larger than zero following a restart indicates that the broker will attempt to store previously consumed messages, as a result causing connection issues.

To resolve this issue, we recommend resolving your XA transactions with either a `rollback()` or a `commit()`. For more information, and to see a code example of resolving XA transactions using `rollback()`, see recovering XA transactions.

I see some of my clients connecting to the broker, while others are unable to connect.

If your broker is in the `RUNNING` status and some clients are able to connect to the broker successfully, while others are unable to do so, you may have reached the wire-level connections limit for the broker. To verify that you've reached the wire-level connections limit, do the following:

- Check the *general* broker logs for your Amazon MQ for ActiveMQ broker in CloudWatch Logs. If the limit has been reached, you will see `Reached Maximum Connections` in the broker logs. For more information on CloudWatch Logs for Amazon MQ for ActiveMQ brokers, see the section called "Understanding the structure of logging in CloudWatch Logs".

Once the wire-level connections limit is reached, the broker will actively refuse additional incoming connections. To resolve this issue, we recommend upgrading your broker instance type. For more information on choosing the best instance type for your workload, see Broker instance types.

If you've confirmed that the number of your wire-level connections is less than the broker connection limit, the issue might be related to rebooting clients. Check your broker logs for numerous and frequent entries of `... Inactive for longer than 600000 ms - removing ...`. The log entry is indicative of rebooting clients or connectivity issues. This effect is more evident when clients connect to the broker via a Network Load Balancer (NLB) with clients that frequently disconnect and reconnect to the broker. This is more typically observed in container based clients.

Check your client-side logs for further details. The broker will clean up inactive TCP connections after 600000 ms, and free up the connection socket.

I'm seeing exception `org.apache.jasper.JasperException: An exception occurred processing JSP page` on the ActiveMQ console when performing operations.

If you are using simple authentication and configuring `AuthorizationPlugin` for queue and topic authorization, make sure to use the `AuthorizationEntries` element in your XML configuration file, and allow the `activemq-webconsole` group permission to all queues and topics. This ensures that the ActiveMQ web console can communicate with the ActiveMQ broker.

The following example `AuthorizationEntry` grants read and write permissions for all queues and topics to the `activemq-webconsole` group.

```
<authorizationEntries>
    <authorizationEntry admin="activemq-webconsole,admins,users" topic=">"
 read="activemq-webconsole,admins,users" write="activemq-webconsole,admins,users" />
    <authorizationEntry admin="activemq-webconsole,admins,users" queue=">"
 read="activemq-webconsole,admins,users" write="activemq-webconsole,admins,users" />
</authorizationEntries>
```

Similarly when integrating your broker with LDAP, make sure to grant permission for the `amazonmq-console-admins` group. For more information about LDAP integration, see the section called “How LDAP integration works”.

Troubleshooting: Amazon MQ for RabbitMQ

Use the information in this section to help you diagnose and resolve common issues you might encounter when working with Amazon MQ for RabbitMQ brokers.

Contents

I can’t see metrics for my queues or virtual hosts in CloudWatch.

If you’re unable to view metrics for your queues or virtual hosts in CloudWatch, check if your queue or virtual host names contain any blank spaces, tabs, or other non-ASCII characters.

Amazon MQ cannot publish metrics for virtual hosts and queues with names containing blank spaces, tabs or other non-ASCII characters.

For more information about dimension names, see Dimension in the *Amazon CloudWatch API Reference*.

How do I enable plugins in Amazon MQ for RabbitMQ?

Amazon MQ for RabbitMQ currently only supports the RabbitMQ management, shovel, federation, consistent-hash exchange plugin, which are enabled by default. For more information on using supported plugins, see the section called "Plugins".

I'm unable to change Amazon VPC configuration for the broker.

Amazon MQ does not support changing Amazon VPC configuration after your broker is created. Please note that you will need to create a new broker with the new Amazon VPC configuration and update the client connection URL with the new broker connection URL.

Troubleshooting: Amazon MQ action required codes

Amazon MQ returns an exception for certain API operations, such as `RebootBroker`, if your broker is in an unhealthy state and requires a set of actions to return to a healthy state. The exceptions include specific *action required codes* that help you to identify a root cause, and address the issue and recover your broker.

Use the following list of topics to identify the action required code you have received, and learn more about the steps we recommend to resolve your issue.

Action required codes

- Amazon MQ for RabbitMQ: High memory alarm
- Amazon MQ for RabbitMQ: Invalid AWS Key Management Service Key
- Amazon MQ for ActiveMQ: Deleted Elastic Network Interface alarm
- Amazon MQ for ActiveMQ: Broker Out Of Memory alarm
- Amazon MQ for RabbitMQ: Disk limit alarm

Amazon MQ for RabbitMQ: High memory alarm

RabbitMQ will raise a high memory alarm when the broker's memory usage, identified by CloudWatch metric `RabbitMQMemUsed`, exceeds the memory limit, identified by `RabbitMQMemLimit`. `RabbitMQMemLimit` is set by Amazon MQ and has been specifically tuned considering the memory available for each host instance type.

An Amazon MQ for RabbitMQ broker that has raised a high memory alarm will block all clients that are publishing messages. Due to high memory usage, your broker might also experience other issues that complicate diagnosis and resolution of the alarm.

Single-instance brokers that can't complete start-up due to high memory usage might enter a restart loop, during which, interactions with the broker are limited. In cluster deployments, queues might experience paused synchronization of messages between replicas on different nodes. Paused queue syncs prevent consumption of messages from queues and must be addressed separately while resolving the memory alarm.

Amazon MQ will not restart a broker experiencing a high memory alarm and will return an exception for `RebootBroker` API operations as long as the broker continues to raise the alarm.

Use the information in this section to help you diagnose and resolve RabbitMQ high memory alarms raised by your broker.

> **Note**
>
> It may take up to several hours for the RABBITMQ_MEMORY_ALARM status to clear after you take the required actions.

> **Note**
>
> You cannot downgrade a broker from an `mq.m5.` instance type to an `mq.t3.micro` instance type. If you wish to downgrade, you must delete your broker and create a new one.

Topics

- Diagnosing high memory alarm using the RabbitMQ web console
- Diagnosing high memory alarm using Amazon MQ metrics

- Addressing high memory alarm
- Reducing the number of connections and channels
- Addressing paused queue synchronizations in cluster deployments
- Addressing restart loops in single-instance brokers
- Preventing high memory alarms

Diagnosing high memory alarm using the RabbitMQ web console

The RabbitMQ web console can generate and display detailed memory usage information for each node. You can find this information by doing the following:

1. Sign in to AWS Management Console and open your broker's RabbitMQ web console.
2. On the RabbitMQ console, on the **Overview** page, choose the name of a node from the **Nodes** list.
3. On the node detail page, choose **Memory details** to expand the section to view the node's memory usage information.

The memory usage information that RabbitMQ provides in the web console can help you determine which resources might be consuming too much memory and contributing to the high memory alarm. For more information about the memory usage details available via the RabbitMQ web console, see Reasoning About Memory Use on the RabbitMQ Server Documentation website.

Diagnosing high memory alarm using Amazon MQ metrics

Amazon MQ enables metrics for your broker by default. You can view your broker metrics by accessing the CloudWatch console, or by using the CloudWatch API. The following metrics are useful when diagnosing the RabbitMQ high memory alarm.

Amazon MQ CloudWatch metric	Reason for high memory use	
`MessageCount`	Messages are stored in memory until they are consumed or discarded. A high message count might indicate overutilization of	

Amazon MQ CloudWatch metric	Reason for high memory use
	resources and can lead to a high memory alarm.
`QueueCount`	Queues are stored in memory, and a high number of queues can lead to a high memory alarm.
`ConnectionCount`	Client connections utilize memory, and too many simultaneous connections can lead to a high memory alarm.
`ChannelCount`	Similar to connections, channels established using each connection are also stored in node memory, and a high number of channels can lead to a high memory alarm.
`ConsumerCount`	For every consumer connected to the broker, a set number of messages are loaded from storage into memory before they are delivered to the consumer. A large number of consumer connections might cause high memory usage and lead to a high memory alarm.

Amazon MQ CloudWatch metric	Reason for high memory use	
`PublishRate`	Publishing messages utilizes the broker' memory. If the rate at which messages are published to the broker is too high and significantly outpaces the rate at which the broker delivers messages to consumers, the broker might raise a high memory alarm.	

Addressing high memory alarm

For each contributor that you identify, we recommended the following set of actions to mitigate and resolve the broker's high memory alarm.

Reason for high memory use	Amazon MQ recommendation	
The **number of messages** in the queues is too high.	Do any of the following: • Consume messages published to the queues. • Purge messages from queues. • Delete the queues from your broker.	
The **number of queues** configured on the broker is too high.	Reduce the number of queues.	

Reason for high memory use	Amazon MQ recommendation
The **number of connections** established on the broker is too high.	Reduce the number of connections. For more information, see the section called "Reducing the number of connections and channels".
The number of channels established on the broker is too high.	Reduce the number of channels. For more information, see the section called "Reducing the number of connections and channels".
The **number of consumers** connected to the broker is too high.	Reduce the number of consumers connected to the broker.
The **message publishing rate** is too high.	Reduce the rate at which publishers send messages to the broker.
The **client connection attempt** rate is too high.	Reduce the frequency at which clients attempt to connect to the broker in order to publish or consume messages, or configure the broker.

Reducing the number of connections and channels

Connections to your Amazon MQ for RabbitMQ broker can be closed either by your client applications, or by manually closing them using the RabbitMQ web console. To close a connection using the RabbitMQ web console do the following.

1. Sign in to AWS Management Console and open your broker's RabbitMQ web console.

2. On the RabbitMQ console, choose the **Connections** tab.
3. On the **Connections** page, under **All connections**, choose the name of the connection you want to close from the list.
4. On the connection details page, choose **Close this connection** to expand the section, then choose **Force Close**. Optionally, you can replace the default text for **Reason** with your own description. Amazon MQ for RabbitMQ will return the reason you specify to the client when you close the connection.
5. Choose **OK** on the dialog box to confirm and close the connection.

When you close a connection, any channels associated with closed connection will also be closed.

> **Note**
>
> Your client applications may be configured to automatically re-establish connections to the broker after they are closed. In this case, closing connections from the broker web console will not be sufficient for reducing connection or channel counts.

For brokers without public access, you can temporarily block connections by denying inbound traffic on the appropriate message protocol port, for example, port 5671 for AMQP connections. You can block the port in the security group that you provided to Amazon MQ when creating the broker. For more information on modifying your security group, see Adding rules to a security group in the *Amazon VPC User Guide*.

Addressing paused queue synchronizations in cluster deployments

While addressing RabbitMQ's high memory alarms, you may find that messages on one or multiple queues cannot be consumed. These queues may be in the process of synchronizing messages between nodes, during which the respective queues become unavailable for publishing and consuming. Queue synchronizations might become paused due to the high memory alarm, and even contribute to the memory alarm.

For information about stopping and retrying paused queue syncs, see the section called “Resolving paused queue sync”.

Addressing restart loops in single-instance brokers

An Amazon MQ for RabbitMQ single-instance broker that raises a high memory alarm is at risk of becoming unavailable if it restarts and does not have enough memory to start up. This can cause RabbitMQ to enter a restart loop and prevent any further interactions with the broker until the issue is resolved. If your broker is in a restart loop, you will not be able to apply the Amazon MQ recommended actions previously described in this section to resolve the high memory alarm.

To recover your broker, we recommend upgrading to a larger instance type with more memory. Unlike in cluster deployments, you can upgrade a single-instance broker while it is experiencing a high memory alarm because there are no queue synchronizations to perform between nodes during a restart.

Preventing high memory alarms

For each contributing factor you identify, we recommends the following set of actions for preventing and reducing the occurrence of RabbitMQ high memory alarms.

Reason high memory use	Amazon MQ recommendation
The **number of messages** in the queues is too high.	Do the following: • Enable lazy queues. • Set, or reduce the queue depth limit.
The **number of queues** configured on the broker is too high.	Set, or reduce the queue count limit.
The **number of connections** established on the broker is too high.	Set, or reduce the connection count limit.
The **number of channels** established on the broker is too high.	Set a maximum number of channels per connection on client applications.

Reason high memory use	Amazon MQ recommendation	
The **number of consumers** connected to the broker is too high.	Set a small consumer pre-fetch limit.	
The **client connection attempt** rate is too high.	Use longer-lived connections to reduce the number and frequency of connection attempts.	

After your broker's memory alarm has been resolved, you can upgrade your host instance type to an instance with additional resources. For information on how to update your broker's instance type see `UpdateBrokerInput` in the *Amazon MQ REST API Reference*.

For a complete listing of broker instance types, see the section called "Amazon MQ for RabbitMQ instance types".

Amazon MQ for RabbitMQ: Invalid AWS Key Management Service Key

Amazon MQ for RabbitMQ will raise an INVALID_KMS_KEY critical action required code when a broker created with a customer managed AWS KMS key(CMK) detects that the AWS Key Management Service (KMS) key is disabled. A RabbitMQ broker with a CMK periodically verifies that the KMS key is enabled and the broker has all necessary grants. If RabbitMQ cannot verify that the key is enabled, the broker is quarantined and RabbitMQ will return INVALID_KMS_KEY.

Without an active KMS key, the broker does not have basic permissions for customer managed KMS keys. The broker cannot perform cryptographic operations using your key until you re-enable your key and the broker restarts. A RabbitMQ broker with a disabled KMS key is quarantined to prevent deterioration. After RabbitMQ determines the KMS key is active again, your broker is removed from quarantine. Amazon MQ does not restart a broker with a disabled KMS key and returns an exception for `RebootBroker` API operations as long as the broker continues to have an invalid KMS key.

Diagnosing and addressing INVALID_KMS_KEY

To diagnose and address the INVALID_KMS_KEY action required code, you must use the AWS Command Line Interface (CLI) and the AWS Key Management Service console.

To re-enable your KMS key

1. Call the `DescribeBroker` method to retrieve the `kmsKeyId` for your CMK broker.
2. Sign in to the AWS Key Management Service console.
3. On the **Customer managed keys** page, locate the KMS Key ID of the problematic broker and verify the status is **Enabled**.
4. If your KMS key has been disabled, re-enable the key by choosing **Key Actions**, then choose **Enable**. After your key has been re-enabled, you must wait for RabbitMQ to remove the broker from quarantine.

To verify that the necessary grants are still associated with the broker's KMS key, call the `ListGrantListGrant` method to verify that `mq_rabbit_grant` and `mq_grant` are present. If the KMS grant or key has been deleted, you must delete the broker and create a new one with all necessary grants. For steps on deleting a broker, see Deleting a broker.

To prevent the INVALID_KMS_KEY critical action required code, do not manually delete or disable a KMS key or CMK grant. If you wish to delete the key, delete the broker first.

Amazon MQ for ActiveMQ: Deleted Elastic Network Interface alarm

Amazon MQ for ActiveMQ will raise a BROKER_ENI_DELETED alarm when you delete a broker's Elastic Network Interface (ENI). When you first create an Amazon MQ broker, Amazon MQ provisions an elastic network interface in the Virtual Private Cloud (VPC) under your account and, thus, requires a number of EC2 permissions.

You must not modify or delete this network interface. Modifying or deleting the network interface can cause a permanent loss of connection between your VPC and your broker. If you wish to delete the network interface, you must delete the broker first.

Amazon MQ for ActiveMQ: Broker Out Of Memory alarm

Amazon MQ for ActiveMQ will raise a BROKER_OOM alarm when the broker undergoes a restart loop due to the insufficient memory capacity. When a broker is in a restart loop, also called a

bounce loop, the broker initiates repeated recovery attempts within a short time window. Brokers that cannot complete start-up due to insufficient memory capacity can enter a restart loop, during which interactions with the broker are limited.

Amazon MQ enables metrics for your broker by default. You can view your broker metrics by accessing the Amazon CloudWatch console, or by using the CloudWatch API. The following metrics are useful when diagnosing the ActiveMQ BROKER_OOM alarm:

Amazon MQ CloudWatch metric	Reason for high memory use
`TotalMessageCount`	Messages are stored in memory until they are consumed or discarded. A high message count might indicate overutilization of resources and can lead to a high memory alarm.
`HeapUsage`	The percentage of the ActiveMQ JVM memory limit that the broker currently uses. A higher percentage indicates the broker is using significant resources and may lead to an OOM alarm.
`ConnectionCount`	Client connections utilize memory, and too many simultaneous connections can lead to a high memory alarm.
`CpuUtilization`	The percentage of allocated EC2 compute units that the broker currently uses.
`TotalConsumerCount`	For every consumer connected to the broker, a

Amazon MQ CloudWatch metric	Reason for high memory use	
	set number of messages are loaded from storage into memory before they are delivered to the consumer. A large number of consumer connections might cause high memory usage and lead to a high memory alarm.	

To prevent restart loops and avoid the BROKER_OOM alarm, ensure that messages are consumed quickly. You can do this by choosing the most effective broker instance type, and also cleaning your Dead Letter Queue to discard undeliverable or expired messages. You can learn more about ensuring effective performance at Amazon MQ for ActiveMQ best practices.

Amazon MQ for RabbitMQ: Disk limit alarm

Disk limit alarm is an indication that the volume of disk used by a RabbitMQ node has decreased due to a high number of messages not consumed while new messages were added. RabbitMQ will raise a disk limit alarm when the broker's free disk space, identified by Amazon CloudWatch metric `RabbitMQDiskFree`, reaches the disk limit, identified by `RabbitMQDiskFreeLimit`. `RabbitMQDiskFreeLimit` is set by Amazon MQ and has been defined considering the disk space available for each broker instance type.

An Amazon MQ for RabbitMQ broker that has raised a disk limit alarm will become unavailable for new messages being published. When running RabbitMQ in a cluster, the disk alarm is cluster-wide. If one node goes under the limit, all other nodes will block incoming messages. Due to the lack of disk space, your broker might also experience other issues that complicate diagnosis and resolution of the alarm.

Amazon MQ will not restart a broker experiencing a disk alarm and will return an exception for `RebootBroker` API operations as long as the broker continues to raise the alarm.

> **Note**
>
> You cannot downgrade a broker from an `mq.m5` instance type to an `mq.t3.micro` instance type. If you wish to downgrade, you must delete your broker and create a new one.

Diagnosing and addressing disk limit alarm

Amazon MQ enables metrics for your broker by default. You can view your broker metrics by accessing the Amazon CloudWatch console, or by using the CloudWatch API.`MessageCount` is a useful metric when diagnosing the RabbitMQ disk limit alarm. Messages are stored in memory until they are consumed or discarded. A high message count indicates overutilization of disk storage and can lead to a disk alarm.

To diagnose the disk limit alarm, use the Amazon MQ Management Console to:

- Consume messages published to the queues.
- Purge messages from queues.
- Delete the queues from your broker.

> **Note**
>
> It may take up to several hours for the RABBITMQ_DISK_ALARM status to clear after you take the required actions.

To prevent the disk limit alarm from reoccurring, you can upgrade your host instance type to an instance with additional resources. For information on how to update your broker's instance type see `UpdateBrokerInput` in the Amazon MQ REST API Reference.

Related resources

Amazon MQ resources

The following table lists useful resources for working with Amazon MQ.

Resource	Description
Amazon MQ REST API Reference	Descriptions of REST resources, example requests, HTTP methods, schemas, parameter s, and the errors that the service returns.
Amazon MQ in the *AWS CLI Command Reference*	Descriptions of the AWS CLI commands that you can use to work with message brokers.
Amazon MQ in the *AWS CloudFormation User Guide*	The `AWS::Amazon MQ::Broker` resource lets you create Amazon MQ brokers, add configuration changes or modify users for the specified broker, return information about the specified broker, and delete the specified broker. The `AWS::Amazon MQ::Configuration` resource lets you create Amazon MQ configurations, add configuration changes or modify users, and return information about the specified configuration.
Regions and Endpoints	Information about Amazon MQ regions and endpoints
Product Page	The primary web page for information about Amazon MQ.
Discussion Forum	A community-based forum for developers to discuss technical questions related to Amazon MQ.

Resource	Description
AWS Premium Support Information	The primary web page for information about AWS Premium Support, a one-on-one, fast-response support channel to help you build and run applications on AWS infrastructure services

Amazon MQ for ActiveMQ resources

The following table lists useful resources for working with Apache ActiveMQ.

Resource	Description
Apache ActiveMQ Getting Started Guide	The official documentation of Apache ActiveMQ.
ActiveMQ in Action	A guide to Apache ActiveMQ that covers the anatomy of JMS messages, connectors, message persistence, authentication, and authorization.
Cross-Language Clients	A list of programming languages and corresponding Apache ActiveMQ libraries. See also ActiveMQ Client and QpidJMS Client.

Amazon MQ for RabbitMQ resources

The following table lists useful resources for working with RabbitMQ.

Resource	Description
The RabbitMQ Getting Started Guide	The official documentation of RabbitMQ.
RabbitMQ Client Libraries and Developer Tools	A guide to the officially supported client libraries and devloper tools for working with

Resource	Description
	RabbitMQ using a variety of programming languages and platforms.
RabbitMQ Best Practices	CloudAMQP's guide on best practices and recommendations for working with RabbitMQ.

Amazon MQ release notes

The following table lists Amazon MQ feature releases and improvements. For changes to the *Amazon MQ Developer Guide*, see Amazon MQ Document History.

Date	Documentation Update
March 4, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ 3.11.28. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.11.28 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
January 19, 2024	Amazon MQ for RabbitMQ does not support the username "guest", and will delete the default guest account when you create a new broker. Amazon MQ will also periodically delete any customer created account called "guest".
December 15, 2023	Amazon MQ is now available in the Israel (Tel Aviv) Region. For information on available regions, see AWS Regions and Endpoints in the *AWS General Reference guide*.
December 11, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ 3.10.25. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.10.25 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog

Date	Documentation Update
	For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
October 26, 2023	Amazon MQ has released the latest ActiveMQ minor versions 5.15.16, 5.16.7, 5.17.6 with a critical update. We have deprecated the older minor versions of ActiveMQ and will be updating all brokers on any version of 5.15 to 5.15.16, or 5.16 to 5.16.7 and 5.17 to 5.17.6. For more information on updating your ActiveMQ broker, see Managing Amazon MQ for ActiveMQ engine versions.
September 27, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ 3.11.20. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.11.20 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.

Date	Documentation Update
July 27, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ 3.11.16 For more information about the fixes and features in this release, see the following: • RabbitMQ 3.11.16 release notes on the RabbitMQ server GitHub repositor y • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
July 27, 2023	Amazon MQ for RabbitMQ now supports creating and applying configura tions to your RabbitMQ broker. For more information on adding confgurations to your broker, see RabbitMQ Broker Configurations. For more information about this feature, see: • Operator policies • Changes to the operator policies

Date	Documentation Update
June 23, 2023	Amazon MQ now supports ActiveMQ 5.17.3, a new minor engine version release. This release supports the new cross-Region data replication (CRDR) feature from Amazon MQ. For more information, see the following: • To get started with CRDR, see Cross-Region data replication for Amazon MQ for ActiveMQ in the Developer Guide. • ActiveMQ 5.17.3 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Upgrading an Amazon MQ broker engine version • Working with Spring XML configuration files
June 21, 2023	Amazon MQ for ActiveMQ now offers a cross-Region data replication (CRDR) feature that allows for asynchronous message replication from the primary broker in a primary AWS Region to the replica broker in a replica Region. If the primary broker in the primary Region fails, you can promote the replica broker in the secondary Region to primary by initiating a switchover or failover. To get started with CRDR, see Cross-Region data replication for Amazon MQ for ActiveMQ in the Developer Guide.
May 18, 2023	Amazon MQ is now available in the following regions: • Asia Pacific (Melbourne) • Asia Pacific (Hyderabad) • Europe (Spain) • Europe (Zurich) For information on available regions, see AWS Regions and Endpoints in the *AWS General Reference guide*.

Date	Documentation Update
April 14, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.9.27. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.9.27 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
April 14, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.10.20. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.10.20 release notes on the RabbitMQ server GitHub repositor y • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.

Date	Documentation Update
March 31, 2023	Amazon MQ for RabbitMQ has disabled RabbitMQ engine version 3.10.17 The Amazon MQ for RabbitMQ team, and the open source maintainers of RabbitMQ, have identified an issue with the RabbitMQ management console on version 3.10.17. Amazon MQ has retracted this version. To mitigate the impacts of this issue, create new brokers with version 3.10.20 while we work to support a new patch version of RabbitMQ. We recommend activating the auto minor version upgrade option to automatically get the latest bug fixes, security updates and performance enhancements. For more information on available Amazon MQ for RabbitMQ versions, see Amazon MQ for RabbitMQ engine versions.
March 1, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.10.17. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.10.17 release notes on the RabbitMQ server GitHub repositor y • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.

Date	Documentation Update
February 21, 2023	Amazon MQ for RabbitMQ now integrates with AWS Key Management Service (KMS) to offer server-side encryption. You can now select your own customer managed CMK, or use an AWS managed KMS key in your AWS KMS account. For more information, see Encryption at rest. Amazon MQ supports using AWS KMS keys in the following ways. • **Amazon MQ owned KMS key (default)** — The key is owned and managed by Amazon MQ and is not in your account. • **AWS managed KMS key** — The AWS managed KMS key (`aws/mq`) is a KMS key in your account that is created, managed, and used on your behalf by Amazon MQ. • **Select existing customer managed KMS key** — Customer managed KMS keys are created and managed by you in AWS Key Management Service (KMS).
January 13, 2023	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.8.34. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.8.34 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.

Date	Documentation Update
December 15, 2022	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.9.24. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.9.24 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
December 13, 2022	Amazon MQ is now available in the Middle East (UAE) Region. For informati on on available regions, see AWS Regions and Endpoints in the *AWS General Reference guide*.
November 14, 2022	Amazon MQ for RabbitMQ now supports 3.10, a major engine version release. You can now enable classic Queues version 2 (CQv2) on your RabbitMQ queues. Direct updates from 3.8 to 3.10 are not supported. For more information, see the following: • RabbitMQ 3.10.10 release notes • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
November 9, 2022	Amazon MQ now supports ActiveMQ 5.17.2, a new minor engine version release. For more information, see the following: • ActiveMQ 5.17.2 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Upgrading an Amazon MQ broker engine version • Working with Spring XML configuration files

Date	Documentation Update
August 17, 2022	Amazon MQ now supports ActiveMQ 5.17.1, a new major engine version release. For more information, see the following: • ActiveMQ 5.17.1 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Upgrading an Amazon MQ broker engine version • Working with Spring XML configuration files
July 14, 2022	Amazon MQ now supports ActiveMQ 5.16.5, a minor engine version release. For more information, see the following: • ActiveMQ 5.16.5 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files • Upgrading an Amazon MQ broker engine version
May 4, 2022	Amazon MQ adds inclusive language for `networkConnector` element in broker configuration. • Creating and configuring an Amazon MQ network of brokers
April 25, 2022	Amazon MQ This release adds the `CRITICAL_ACTION_REQUIRED` broker state and the `ActionRequired` API property. `CRITICAL_ACTION_REQUIRED` informs you when your broker is degraded. `ActionRequired` provides you with a code which you can use to find instructions in the Developer Guide on how to resolve the issue. • the section called “Troubleshooting: Amazon MQ action required codes” • `ActionRequired` documentation in the *Amazon MQ API Reference*.

Date	Documentation Update
April 20, 2022	Amazon MQ now supports ActiveMQ 5.16.4, a minor engine version release. For more information, see the following: • ActiveMQ 5.16.4 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files • Upgrading an Amazon MQ broker engine version
March 1, 2022	Amazon MQ is now available in the Asia Pacific (Jakarta) Region. For information on available regions, see AWS Regions and Endpoints in the *AWS General Reference guide*.
February 25, 2022	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.8.27. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.8.27 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
February 16, 2022	Amazon MQ is now available in the Africa (Cape Town) Region. For informati on on available regions, see AWS Regions and Endpoints in the *AWS General Reference guide*.

Date	Documentation Update
February 14, 2022	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.9.13. Automatic minor version upgrades cannot be used to upgrade from Rabbit 3.8 to 3.9. To do so, manually upgrade your broker. For more information on new features introduced in RabbitMQ 3.9, see the release notes page for version 3.9.0 on the GitHub website. **Note** Currently, Amazon MQ does not support streams, or using structure d logging in JSON, introduced in RabbitMQ 3.9. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.9.13 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
February 07, 2022	Amazon MQ for RabbitMQ introduces new broker metrics, allowing you to monitor average resource utilization across all three nodes in a cluster deployment. For more information, see the following: • the section called “Logging and monitoring Amazon MQ for RabbitMQ brokers”

Date	Documentation Update
January 18, 2022	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.8.26. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.8.26 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
January 13, 2022	Amazon MQ introduces the `RABBITMQ_MEMORY_ALARM` status code to inform you when your broker has raised a high memory alarm and is in an unhealthy state. Amazon MQ provides detailed information and recommendations to help you diagnose, resolve and prevent high memory alarms. For more information, see the following. • the section called "`RABBITMQ_MEMORY_ALARM`"
January 6, 2022	When you configure CloudWatch Logs for Amazon MQ for ActiveMQ brokers, Amazon MQ supports using the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in IAM resource-based policies to prevent the confused deputy problem. For more information, see the following. • the section called "Cross-service confused deputy prevention"
December 20, 2021	Amazon MQ for ActiveMQ introduces a set of new metrics, allowing you to monitor the maximum number of connections you can make to your broker using different supported transport protocols, as well as an additional new metric that allows you to monitor the number of nodes connected to your broker in a network of brokers. For more information, see the following. • the section called "Logging and monitoring Amazon MQ for ActiveMQ brokers"

Date	Documentation Update
November 16, 2021	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.8.23. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.8.23 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions.
October 12, 2021	Amazon MQ now supports ActiveMQ 5.16.3, a minor engine version release. For more information, see the following: • ActiveMQ 5.16.3 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Upgrading an Amazon MQ broker engine version • Working with Spring XML configuration files
September 8, 2021	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.8.22. This release includes a fix for an issue with queues using per-message TTL (time to live), identified in the previously supported version, RabbitMQ 3.8.17. We recommend upgrading your existing brokers to version 3.8.22. For more information about the fixes and features in this release, see the following: • RabbitMQ 3.8.22 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog For more information about supported Amazon MQ for RabbitMQ versions and broker upgrades, see Managing Amazon MQ for RabbitMQ engine versions

Date	Documentation Update
August 25, 2021	Amazon MQ for RabbitMQ has temporarily disabled RabbitMQ engine version 3.8.17 due to an issue identified with queues using per-message time-to-live (TTL). We recommend using version 3.8.11.
July 29, 2021	Amazon MQ for RabbitMQ now supports RabbitMQ version 3.8.17. For more information about the fixes and features contained in this update, see the following: • RabbitMQ 3.8.17 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog • Managing Amazon MQ for RabbitMQ engine versions
July 16, 2021	You can now adjust the maintenance window of an Amazon MQ broker using the AWS Management Console, AWS CLI, or the Amazon MQ API. To learn more about broker maintenance windows, see the following. • Maintaining an Amazon MQ broker
July 6, 2021	Amazon MQ for RabbitMQ introduces support for the Consistent Hash exchange type. Consistent Hash exchanges route messages to queues based on a hash value calculated from the routing key of a message. For more information, see the following: • Consistent Hash exchange plugin • RabbitMQ Consistent Hash Exchange Type on the RabbitMQ GitHub repository
June 7, 2021	Amazon MQ now supports ActiveMQ 5.16.2, a new major engine version release. For more information, see the following: • ActiveMQ 5.16.2 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Upgrading an Amazon MQ broker engine version • Working with Spring XML configuration files

Date	Documentation Update
May 26, 2021	Amazon MQ for RabbitMQ is now available in the China (Beijing) and China (Ningxia) Regions. For information on available regions, see AWS Regions and Endpoints.
May 18, 2021	Amazon MQ for RabbitMQ implements broker defaults. When you first create a broker, Amazon MQ creates a set of broker policies and vhost limits based on the instance type and deployment mode you choose, in order to optimize the broker's performance. For more informati on, see the following: • Amazon MQ for RabbitMQ broker defaults
May 5, 2021	Amazon MQ now supports ActiveMQ 5.15.15. For more information, see the following: • ActiveMQ 5.15.15 Release Page • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files
May 5, 2021	Amazon MQ started tracking changes to AWS managed policies. For more information, see the following: • the section called “AWS managed policies”
April 14, 2021	Amazon MQ is now available in the China (Beijing) and China (Ningxia) Regions. For information on available regions, see AWS Regions and Endpoints.
April 7, 2021	Amazon MQ now supports RabbitMQ 3.8.11. For more information about the fixes and features contained in this update, see the following: • RabbitMQ 3.8.11 release notes on the RabbitMQ server GitHub repository • RabbitMQ changelog • Managing Amazon MQ for RabbitMQ engine versions

Date	Documentation Update
April 1, 2021	Amazon MQ is now available in the Asia Pacific (Osaka) Region. For informati on about available regions, see Amazon MQ regions and endpoints.
December 21, 2020	Amazon MQ now supports ActiveMQ 5.15.14. For more information, see the following: • ActiveMQ 5.15.14 Release Notes • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files • **Important** Due to a known Apache ActiveMQ issue in this release, the new **Pause Queue** button in the ActiveMQ web console cannot be used with Amazon MQ for ActiveMQ brokers. For more information about this issue, see AMQ-8104.

Date	Documentation Update
November 4, 2020	Amazon MQ now supports RabbitMQ, a popular open source message broker. This enables you to migrate your existing RabbitMQ message brokers to AWS without having to rewrite code. Amazon MQ for RabbitMQ manages both individual and clustered message brokers and handles tasks like provisioning the infrastructure, setting up the broker, and updating the software. • Amazon MQ supports RabbitMQ 3.8.6. For more information about supported engine versions, see the section called "Version management". • The AWS Free Tier includes up to 750 hours of a single-instance `mq.t3.micro` broker and up to 20GB of storage per month for one year. For more information about supported instance types, see Broker instance types. • With Amazon MQ for RabbitMQ, you can access your brokers using AMQP 0-9-1, and with any language supported by the RabbitMQ client libraries . For more information about supported protocols and cipher suites, see the section called "Amazon MQ for RabbitMQ protocols". • Amazon MQ for RabbitMQ is available in all regions that Amazon MQ is currently available. To learn more about all of the available regions, see the AWS Region Table. To get started with using Amazon MQ, create a broker, and connect a JVM-based application to your RabbitMQ broker, see the section called "Creating and connecting to a RabbitMQ broker".
October 22, 2020	Amazon MQ supports ActiveMQ 5.15.13. For more information, see the following: • ActiveMQ 5.15.13 Release Notes • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files

Date	Documentation Update
September 30, 2020	Amazon MQ is now available in the Europe (Milan) Region. For information about available regions, see Amazon MQ regions and endpoints.
July 27, 2020	You can authenticate Amazon MQ users using the credentials stored in your Active Directory or other LDAP server. You can also add, delete, and modify Amazon MQ users and assign permissions to topics and queues. For more information, see Integrate LDAP with ActiveMQ.
July 17, 2020	Amazon MQ now supports the `mq.t3.micro` instance type. For more information, see Broker instance types.
June 30, 2020	Amazon MQ supports ActiveMQ 5.15.12. For more information, see the following: • ActiveMQ 5.15.12 Release Notes • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files
April 30, 2020	Amazon MQ supports a new child collection element, `systemUsage`, on the `broker` element. For more information, see systemUsage. Amazon MQ also supports three new attributes on the `kahaDB` child element. • `journalDiskSyncInterval` - Interval (ms) for when to perform a disk sync if `journalDiskSyncStrategy=periodic`. • `journalDiskSyncStrategy` - configures the disk sync policy. • `preallocationStrategy` - configures how the broker will try to preallocate the journal files when a new journal file is needed. For more information, see Attributes.

Date	Documentation Update
March 3, 2020	Amazon MQ supports two new CloudWatch metrics • `TempPercentUsage` - The percentage of available temporary storage used by non-persistent messages. • `JobSchedulerStorePercentUsage` - The percentage of disk space used by the job scheduler store. For more information, see Monitoring Amazon MQ using CloudWatch.
February 4, 2020	Amazon MQ is available in the Asia Pacific (Hong Kong) and Middle East (Bahrain) regions. For information on available regions, see AWS Regions and Endpoints.
January 22, 2020	Amazon MQ supports ActiveMQ 5.15.10. For more information, see the following: • ActiveMQ 5.15.10 Release Notes • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files
December 19, 2019	Amazon MQ is available in the Europe (Stockholm) and South America (São Paulo) regions. For information on available regions, see AWS Regions and Endpoints.

Date	Documentation Update
December 16, 2019	Amazon MQ supports creating throughput-optimized brokers by using Amazon Elastic Block Store (EBS)—instead of the default Amazon Elastic File System (Amazon EFS)—for broker storage. To take advantage of high durability and replication across multiple Availability Zones, use Amazon EFS. To take advantage of low latency and high throughput, use Amazon EBS. **Important** • You can use Amazon EBS only with the `mq.m5` broker instance type family. • Although you can change the *broker instance type*, you can't change the *broker storage type* after you create the broker. • Amazon EBS replicates data within a single Availability Zone and doesn't support the ActiveMQ active/standby deployment mode. For more information, see the following: • Storage • Choose the correct broker storage type for the best throughput • The `storageType` property of the `broker-instance-options` resource in the *Amazon MQ REST API Reference* • The `BurstBalance`, `VolumeReadOps`, and `VolumeWriteOps` metrics in the Amazon MQ for ActiveMQ metrics section.
October 18, 2019	Two Amazon CloudWatch metrics are available: `TotalEnqueueCount` and `TotalDequeueCount`. For more information, see ActiveMQ destination (queue and topic) metrics.

Date	Documentation Update
October 11, 2019	Amazon MQ now supports Federal Information Processing Standard 140-2 (FIPS) compliant endpoints in U.S. commercial regions. For more information see the following: • Federal Information Processing Standard (FIPS) 140-2 • Amazon MQ Regions and Endpoints
September 30, 2019	Amazon MQ now includes the ability to scale your brokers by changing the host instance type. For more information, see the `hostInstanceType` property of `UpdateBrokerInput`, and the `pendingHostInstanceType` property of `DescribeBrokerOutput`.
August 30, 2019	You can now update the security groups associated with a broker, both in the console and with `UpdateBrokerInput`.
July 22, 2019	Amazon MQ integrates with AWS Key Management Service (KMS) to offer server-side encryption. You can now select your own customer managed CMK, or use an AWS managed KMS key in your AWS KMS account. For more information, see Encryption at rest. Amazon MQ supports using AWS KMS keys in the following ways. • **AWS owned KMS key** — The key is owned Amazon MQ and is not in your account. • **AWS managed KMS key** — The AWS managed KMS key (aws/mq) is a KMS key in your account that is created, managed, and used on your behalf by Amazon MQ. • **Select existing customer managed CMK** — Customer managed CMKs are created and managed by you in AWS Key Management Service (KMS).
June 19, 2019	Amazon MQ is available in the Europe (Paris) and Asia Pacific (Mumbai) regions. For information on available regions, see AWS Regions and Endpoints.

Date	Documentation Update
June 12, 2019	Amazon MQ is available in the Canada (Central) region. For information on available regions, see AWS Regions and Endpoints.
June 3, 2019	Two new Amazon CloudWatch metrics are available: `EstablishedConnectionsCount` and `InactiveDurableSubscribers`. For more information, see the following: • Monitoring Amazon MQ using CloudWatch • Amazon MQ for ActiveMQ metrics
May 10, 2019	Data storage for new `mq.t2.micro` instance types is limited to 20 GB. For more information, see the following: • the section called "Data Storage" • Broker instance types
April 29, 2019	You can now use tag-based policies and resource-level permissions. For more information, see the following: • How Amazon MQ works with IAM • Resource-level permissions for Amazon MQ API actions
April 16, 2019	You can now retrieve information about broker engine and broker instance options using the REST API. For more information, see the following: • Broker instance options • Broker engine types
April 8, 2019	Amazon MQ supports ActiveMQ 5.15.9. For more information, see the following: • ActiveMQ 5.15.9 Release Notes • Managing Amazon MQ for ActiveMQ engine versions • Working with Spring XML configuration files

Date	Documentation Update
March 4, 2019	Improved the documentation for configuring dynamic failover and the rebalancing of clients for a network of brokers. Enable dynamic failover by configuring `transportConnectors` along with `networkConnectors` configuration options. For more information, see the following: • Dynamic Failover With Transport Connectors • Amazon MQ Network of brokers • Amazon MQ Broker Configuration Parameters
February 27, 2019	Amazon MQ is available in the Europe (London) Region in addition to the following regions: • Asia Pacific (Singapore) • US East (Ohio) • US East (N. Virginia) • US West (N. California) • US West (Oregon) • Asia Pacific (Tokyo) • Asia Pacific (Seoul) • Asia Pacific (Sydney) • Europe (Frankfurt) • Europe (Ireland)
January 24, 2019	The default configuration now includes a policy to purge inactive destinati ons.
January 17, 2019	Amazon MQ `mq.t2.micro` instance types now support only 100 connectio ns per wire-level protocol. For more information, see, Quotas in Amazon MQ.

Date	Documentation Update
December 19, 2018	You can configure a series of Amazon MQ brokers in a network of brokers. For more information, see the following sections: • Amazon MQ Network of brokers • Creating and Configuring a Network of Brokers • Configure Your Network of Brokers Correctly • `networkConnector` • `networkConnectionStartAsync`

December 11, 2018	Amazon MQ supports ActiveMQ 5.15.8, 5.15.6, and 5.15.0.
	• Resolved bugs and improvements in ActiveMQ:
	  • ActiveMQ 5.15.8 Release Notes
	  • ActiveMQ 5.15.7 Release Notes

December 5, 2018	AWS supports resource tagging to help track your cost allocation. You can tag resources when creating them, or by viewing the details of that resource. For more information, see Tagging resources.
November 19, 2018	AWS has expanded its SOC compliance program to include Amazon MQ as an SOC compliant service.
October 15, 2018	• The maximum number of groups per user is 20. For more information, see Users. • The maximum number of connections per broker, per wire-level protocol is 1,000. For more information, see Brokers.
October 2, 2018	AWS has expanded its HIPAA compliance program to include Amazon MQ as a HIPAA Eligible Service.

Date	Documentation Update

September 27, 2018	Amazon MQ supports ActiveMQ 5.15.6, in addition to 5.15.0. For more information, see the following:
	• Editing broker engine version, instance type, CloudWatch logs, and maintenance preferences
	• Resolved bugs and improvements in the ActiveMQ documentation:
	  • ActiveMQ 5.15.6 Release Notes
	  • ActiveMQ 5.15.5 Release Notes
	  • ActiveMQ 5.15.4 Release Notes
	  • ActiveMQ 5.15.3 Release Notes
	  • ActiveMQ 5.15.2 Release Notes
	  • ActiveMQ 5.15.1 Release Notes
	• ActiveMQ Client 5.15.6

August 31, 2018	• The following metrics are available:
	  • `CurrentConnectionsCount`
	  • `TotalConsumerCount`
	  • `TotalProducerCount`
	  For more information, see the Amazon MQ for ActiveMQ metrics section.
	• The IP address of the broker is displayed on the **Details** page.
	  Note
	  For brokers with public accessibility disabled, the internal IP address is displayed.

Date	Documentation Update
August 30, 2018	Amazon MQ is available in the Asia Pacific (Singapore) Region in addition to the following regions: • US East (Ohio) • US East (N. Virginia) • US West (N. California) • US West (Oregon) • Asia Pacific (Tokyo) • Asia Pacific (Seoul) • Asia Pacific (Sydney) • Europe (Frankfurt) • Europe (Ireland)
July 30, 2018	You can configure Amazon MQ to publish general and audit logs to Amazon CloudWatch Logs. For more information, see Configuring Amazon MQ to publish logs to Amazon CloudWatch Logs.
July 25, 2018	Amazon MQ is available in the Asia Pacific (Tokyo) and Asia Pacific (Seoul) Regions in addition to the following regions: • US East (Ohio) • US East (N. Virginia) • US West (N. California) • US West (Oregon) • Asia Pacific (Sydney) • Europe (Frankfurt) • Europe (Ireland)
July 19, 2018	You can use AWS CloudTrail to log Amazon MQ API calls. For more informati on, see Logging Amazon MQ API calls using CloudTrail.

Date	Documentation Update
June 29, 2018	In addition to `mq.t2.micro` and `mq.m4.large`, the following broker instance types are available for regular development, testing, and production workloads that require high throughput: • `mq.m5.large` • `mq.m5.xlarge` • `mq.m5.2xlarge` • `mq.m5.4xlarge` For more information, see Broker instance types.
June 27, 2018	Amazon MQ is available in the US West (N. California) Region in addition to the following regions: • US East (Ohio) • US East (N. Virginia) • US West (Oregon) • Asia Pacific (Sydney) • Europe (Frankfurt) • Europe (Ireland)

Date	Documentation Update

June 14, 2018	• You can use the `AWS::Amazon MQ::Broker` AWS CloudFormation resource to perform the following actions:
	  • Create a broker.
	  • Add configuration changes or modify users for the specified broker.
	  • Return information about the specified broker.
	  • Delete the specified broker.
	
	  Note
	  When you change any property of the Amazon MQ Broker ConfigurationId or Amazon MQ Broker User property type, the broker is rebooted immediately.
	
	• You can use the `AWS::Amazon MQ::Configuration` AWS CloudFormation resource to perform the following actions:
	  • Create a configuration.
	  • Update the specified configuration.
	  • Return information about the specified configuration.
	
	  Note
	  You can use AWS CloudFormation to modify—but not delete—an Amazon MQ configuration.

June 7, 2018	The Amazon MQ console supports German, Brazilian Portuguese, Spanish, Italian, and Traditional Chinese.
May 17, 2018	The limit of number of users per broker is 250. For more information, see Users.
March 13, 2018	Creating a broker takes about 15 minutes. For more information, see Finish creating the broker.

Date	Documentation Update
March 1, 2018	• You can configure the concurrent store and dispatch for Apache KahaDB using the `concurrentStoreAndDispatchQueues` attribute. • The `CpuCreditBalance` CloudWatch metric is available for `mq.t2.micro` broker instance type.
January 10, 2018	The following changes affect the Amazon MQ console: • In the broker list, the **Creation** column is hidden by default. To customize the page size and columns, choose . • On the ***MyBroker*** page, in the **Connections** section, choosing the name of your security group or opens the EC2 console (instead of the VPC console). The EC2 console allows more intuitive configuration of inbound and outbound rules. For more information, see the updated Enable inbound connections section.
January 9, 2018	• The permission for REST operation ID `UpdateBroker` is listed correctly as `mq:UpdateBroker` on the IAM console. • The erroneous `mq:DescribeEngine` permission is removed from the IAM console.

Date	Documentation Update

November 28, 2017	This is the initial release of Amazon MQ and the *Amazon MQ Developer Guide*.
	
	• Amazon MQ is avaialble in the following regions:
	  • US East (Ohio)
	  • US East (N. Virginia)
	  • US West (Oregon)
	  • Asia Pacific (Sydney)
	  • Europe (Frankfurt)
	  • Europe (Ireland)
	
	  Using the `mq.t2.micro` instance type is subject to *CPU credits and baseline performance*—with the ability to *burst* above the baseline level (for more information, see the `CpuCreditBalance` metric). If your application requires *fixed performance*, consider using an `mq.m5.large` instance type.
	• You can create `mq.m4.large` and `mq.t2.micro` brokers.
	
	  Using the `mq.t2.micro` instance type is subject to *CPU credits and baseline performance*—with the ability to *burst* above the baseline level (for more information, see the `CpuCreditBalance` metric). If your application requires *fixed performance*, consider using an `mq.m5.large` instance type.
	• You can use the ActiveMQ 5.15.0 broker engine.
	• You can also create and manage brokers programmatically using Amazon MQ REST API and AWS SDKs.
	• You can access your brokers by using any programming language that ActiveMQ supports and by enabling TLS explicitly for the following protocols:
	  • AMQP
	  • MQTT
	  • MQTT over WebSocket
	  • OpenWire

Date	Documentation Update
	• STOMP • STOMP over WebSocket • You can connect to ActiveMQ brokers using various ActiveMQ clients. We recommend using the ActiveMQ Client. For more information, see Connecting a Java application to your broker. • Your broker can send and receive messages of any size.

Amazon MQ Document History

The following table lists changes to the *Amazon MQ Developer Guide*. For Amazon MQ feature releases and improvements, see Amazon MQ release notes.

Date	Documentation Update
August 22, 2022	Created separate parent chapters for Amazon MQ for ActiveMQ and Amazon MQ for RabbitMQ engines. These parent chapters now contain engine details, tutorials, and best practices: • Working with Amazon MQ for ActiveMQ • Working with Amazon MQ for RabbitMQ
January 13, 2022	Added a new troubleshooting section that lists the status codes that Amazon MQ returns when a broker is in an unhealthy state, along with detailed information about diagnosing, and recovering the broker. • the section called “Troubleshooting: Amazon MQ action required codes”
November 8, 2021	Added new tutorial that describes setting up a Python Pika client with Amazon MQ for RabbitMQ brokers. • the section called “Using Python Pika with Amazon MQ for RabbitMQ”
October 8, 2021	Added the following troubleshooting topics for both Amazon MQ for ActiveMQ and Amazon MQ for RabbitMQ broker engines:

Date	Documentation Update
	• Some clients unable to connect • the section called "How do I enable plugins in Amazon MQ for RabbitMQ?" • the section called "I'm unable to change Amazon VPC configuration for the broker."
September 22, 2021	Added the following topics for troubleshooting common connection, and authorization issues with Amazon MQ for ActiveMQ brokers: • Connecting to broker after a restart • JSP exception on the web console
August 12, 2021	Added the following section to describe troubleshooting common issues when working with Amazon MQ brokers. • *Troubleshooting*
July 29, 2021	Added the following sections to describe Amazon MQ for RabbitMQ version management and upgrading Amazon MQ brokers to new minor and major engine versions as they are supported. • the section called "Version management"
July 21, 2021	Added the following sections to describe connecting an Amazon MQ broker to AWS Lambda as an event source. • Connect your Amazon MQ for ActiveMQ broker to Lambda • Connect your Amazon MQ for RabbitMQ broker to Lambda
July 16, 2021	Added the following sections to describe Amazon MQ broker maintenance windows, and how to adjust a maintenance window using the AWS Management Console, AWS CLI, or the Amazon MQ API. • the section called "Maintaining a broker"

Date	Documentation Update
June 7, 2021	Added the following sections to describe Amazon MQ for ActiveMQ version management and upgrading Amazon MQ brokers to new minor and major engine versions as they are supported. • the section called "Version management" • the section called "Upgrading the engine version"
May 18, 2021	Added the following section to describe Amazon MQ for RabbitMQ broker defaults • the section called "Broker defaults"
May 5, 2021	Added the following section for describing AWS managed policies for Amazon MQ and updates to these policies: • the section called "AWS managed policies"
February 16, 2021	Added the following tutorial section for Amazon MQ for RabbitMQ: • the section called "Resolving paused queue sync"

November 4, 2020	• Added the following sections to document Amazon MQ for RabbitMQ support:
	  • the section called "Creating and connecting to a RabbitMQ broker"
	  • the section called "RabbitMQ tutorials"
	  • the section called "Amazon MQ for RabbitMQ best practices"
	  • the section called "RabbitMQ engine"
	  • the section called "Configuring Amazon MQ for RabbitMQ logs"
	  • the section called "Using service-linked roles"
	• Additional revisions to existing chapters and sections of the guide were made to accurately document Amazon MQ for RabbitMQ support.

Date	Documentation Update

December 16, 2019	• Added the following sections:
	  • Storage
	  • Choose the correct broker storage type for the best throughput
	• Revised the information in the following sections:
	  • Broker
	  • Broker instance types
	  • Amazon MQ single-instance broker
	  • Amazon MQ active/standby broker for high availability
	  • Create an ActiveMQ broker
	  • Creating and configuring a broker

July 19, 2019	Modified and added content on encryption management in the following sections:
	• Data protection in Amazon MQ
	  • Encryption at rest
	  • Encryption in transit
	• `EncryptionOptions`

April 22, 2019	Added the following sections for tag-based policies and resource-level permissions: • How Amazon MQ works with IAM • Resource-level permissions for Amazon MQ API actions
March 4, 2019	Improved the documentation for configuring dynamic failover and the rebalancing of clients for a network of brokers. Enable dynamic failover by configuring `transportConnectors` along with `networkConnectors` configuration options. • Dynamic Failover With Transport Connectors • Amazon MQ Network of brokers • Amazon MQ Broker Configuration Parameters

Date	Documentation Update
January 5, 2019	Improved documentation on some per-minute metrics. For more informati on, see the following: ActiveMQ destination (queue and topic) metrics.

December 19, 2018	• Added the following sections:
	  • Amazon MQ Network of brokers
	  • Creating and Configuring a Network of Brokers
	  • Configure Your Network of Brokers Correctly
	  • `networkConnector`
	  • `networkConnectionStartAsync`
	• Added the `networkConnectors` child collection element to the Elements, Child Collection Elements, and Their Child Elements Permitted in Amazon MQ Configurations section.

December 11, 2018	Updated documentation to reflect availability of ActiveMQ version 5.15.8.
December 5, 2018	Added the Tagging resources section.
October 26, 2018	Added the Avoid slow restarts by recovering prepared XA transactions section.
October 15, 2018	Updated the Quotas in Amazon MQ section.
October 1, 2018	Corrected the information in the Next steps section.

September 27, 2018	• Added the Editing broker engine version, instance type, CloudWatch logs, and maintenance preferences section.
	• Updated the following sections:
	  • Create an ActiveMQ broker
	  • Configure Basic Broker Settings

Date	Documentation Update
September 18, 2018	Added the following note to the Creating and managing ActiveMQ broker users section: You can't configure groups independently of users. A group label is created when you add at least one user to it and deleted when you remove all users from it.
August 31, 2018	• Clarified the terminology for active/standby brokers. For more informati on, see Amazon MQ active/standby broker for high availability. • Simplified the terminology for the maintenance window. For more information, see Amazon MQ broker configuration lifecycle. • Rewrote the Configure Additional Broker Settings section. • Updated the Amazon MQ for ActiveMQ metrics and Listing brokers and viewing broker details sections.
August 15, 2018	Corrected the information in the Create an ActiveMQ broker section.
August 13, 2018	Added the Accessing the broker web console without public accessibility section.
August 2, 2018	• Added the Troubleshooting CloudWatch Logs Configuration section. • Added the following admonition throughout this guide: **Important** In the following example code, producers and consumers run in a single thread. For production systems (or to test broker instance failover), make sure that your producers and consumers run on separate hosts or threads.
August 1, 2018	Corrected the information in the following sections: • Understanding the structure of logging in CloudWatch Logs • Connect a Java application to your broker

Date	Documentation Update
July 31, 2018	• Moved the 3-minute demo video to the Getting Started with Amazon MQ section. • Added the 3-minute getting started video to the What is Amazon MQ? section.
July 30, 2018	• Added the Configuring Amazon MQ to publish logs to Amazon CloudWatch Logs section. • Updated the Configure Additional Broker Settings section.
July 19, 2018	• Added the Logging Amazon MQ API calls using CloudTrail section.
July 5, 2018	• Added an `authorizationEntry` child element cross-reference to the Always configure an authorization map section. • Clarified the information in the Integrating ActiveMQ brokers with LDAP section. • Clarified the information in the API Throttling section.
June 29, 2018	• Updated the information in the Broker instance types section. • Added the Choose the Correct Broker Instance Type for the Best Throughput section.
June 4, 2018	In addition to GitHub, HTML, PDF, and Kindle, the *Amazon MQ Developer Guide* release notes are available as an RSS feed. 
May 29, 2018	Made the following changes in the Working Java Example section: • Added a STOMP+WSS Java example. The STOMP+WSS example Java code connects to a broker, creates a queue, and publishes and receives a message. • Improved the MQTT Java example. • Improved the OpenWire Java example.

Date	Documentation Update
May 24, 2018	Corrected the wire-level protocol endpoint port in the MQTT Java example in the Working Java Example section.
May 22, 2018	Corrected the information in all Java dependency sections.
May 17, 2018	Corrected the information in the Users section.
May 15, 2018	Corrected the information in the Ensuring effective Amazon MQ performance section.
May 8, 2018	• Placed the Amazon MQ REST API permissions reference in its own section. • Created the IAM Permissions Required to Create an Amazon MQ Broker section with an example custom IAM policy.
May 7, 2018	• Clarified throughout this guide that the broker maintenance window is 2 hours long. For more information, see Amazon MQ broker configuration lifecycle. • Added explanations for why the `ec2:CreateNetworkInterface` and `ec2:CreateNetworkInterfacePermission` permissions are necessary for creating a broker. For more information, see API authentication and authorization for Amazon MQ.
May 1, 2018	Clarified the information about the maintenance window for active/standby brokers in the following sections: • Amazon MQ active/standby broker for high availability • Creating and configuring a broker • Creating and applying broker configurations

Date	Documentation Update
April 27, 2018	Rewrote the following sections and optimized example Java code to match the recommendation to use connection pooling only for producers, *not* consumers: • Always Use Connection Pooling • Create a message producer and send a message • Create a message consumer and receive the message • AmazonMQExample.java
April 26, 2018	Added an MQTT Java example to the Working Java Example section. The MQTT example Java code connects to a broker, creates a topic, and publishes and receives a message.
April 4, 2018	Renamed the Communicating with Amazon MQ section to Connecting to Amazon MQ.
April 3, 2018	Clarified and corrected the information in the Disable Concurrent Store and Dispatch for Queues with Slow Consumers section.
April 2, 2018	Moved the Concurrent Store and Dispatch for Queues in Amazon MQ section to the Disable Concurrent Store and Dispatch for Queues with Slow Consumers section.

March 27, 2018	• Replaced the re:Invent launch video with a 3-minute demo video in the What is Amazon MQ? section.
	• Restructured the following sections:
	  • Broker Architecture
	  • How Amazon MQ Works.
	• Moved Amazon MQ broker configuration lifecycle under the Broker Architecture section.

March 22, 2018	Clarified the following statement throughout this guide: Amazon MQ encrypts messages at rest and in transit using encryption keys that it manages and stores securely. For more information, see the *AWS Encryption SDK Developer Guide*.

Date	Documentation Update
March 19, 2018	Clarified the following statement throughout this guide: An **Active/standby broker** is comprised of two brokers in two different Availability Zones, configured in a *redundant pair*. These brokers communicate synchronously with your application, and with Amazon EFS.
March 15, 2018	• Restructured the Amazon MQ Basic elements section.
March 12, 2018	• Clarified and corrected the information in the Security best practices for Amazon MQ and Connecting to Amazon MQ sections. • Added the Disable Concurrent Store and Dispatch for Queues with Slow Consumers section. • Grouped admonitions into a preface for the Configure advanced broker settings section.
March 9, 2018	• Clarified and corrected the information in the Always configure an authorization map section. • Added the authorizationEntry section and updated the kahaDB section.
March 8, 2018	• Added the Always configure an authorization map section. • Added notes about broker suffixes to the Monitoring Amazon MQ using CloudWatch section.
March 6, 2018	Added the following note throughout this guide: **Note** Using the `mq.t2.micro` instance type is subject to *CPU credits and baseline performance*—with the ability to *burst* above the baseline level (for more information, see the `CpuCreditBalance` metric). If your application requires *fixed performance*, consider using an `mq.m5.large` instance type.

Date	Documentation Update
March 1, 2018	• Added the `CpuCreditBalance` metric to the Amazon MQ for ActiveMQ metrics section. • Added the Amazon MQ Child Element Attributes section. • Added links from elements in the the section called "Permitted Elements" section to their attributes and to child collection elements. • Made corrections to the AWS Glossary in GitHub.
February 28, 2018	Corrected image display in GitHub.
February 27, 2018	In addition to HTML, PDF, and Kindle, the *Amazon MQ Developer Guide* is available on GitHub. To leave feedback, choose the GitHub icon in the upper right-hand corner. 
February 26, 2018	• Made regions consistent in all examples and diagrams. • Optimized links to the AWS console and product webpages.
February 22, 2018	Clarified and corrected the information in the following sections: • Prefer brokers without public accessibility • Always Use the Failover Transport to Connect to Multiple Broker Endpoints • API authentication and authorization for Amazon MQ • Integrating ActiveMQ brokers with LDAP
February 21, 2018	Corrected the Java code in the following sections: • Working Java Example • Connect a Java application to your broker • Always Use Connection Pooling

Date	Documentation Update
February 20, 2018	Clarified and corrected the information in the Security in Amazon MQ and Best Practices sections.
February 19, 2018	• Corrected the Java code in the Always Use Connection Pooling section. • Restructured and expanded the Best Practices sections and Security in Amazon MQ sections.

February 16, 2018	• Added the Security best practices for Amazon MQ section.
	• Updated the Connecting to Amazon MQ section.
	• Corrected the Java code in the following sections:
	  • Getting Started with Amazon MQ
	  • AmazonMQExample.java

February 15, 2018	• Restructured and expanded the Best Practices sections section.
	• Updated the following sections:
	  • How can I get started with Amazon MQ?
	  • Next steps (Getting Started)
	  • Related resources

February 14, 2018	Updated the following sections: • Quotas in Amazon MQ • API Throttling • Security in Amazon MQ
February 13, 2018	• Updated the Related resources section. • Updated the Quotas in Amazon MQ section. • Added the We want to hear from you section.
January 25, 2018	• Fixed an error in the Add Java dependencies subsection of the Working Java Example section. • The permission for REST operation ID `RebootBroker` is listed correctly as `mq:RebootBroker` on the IAM console.

Date	Documentation Update
January 24, 2018	• Added the Never Modify or Delete the Amazon MQ Elastic Network Interface section. • Updated all diagrams throughout this guide. • Added links to the *Amazon MQ REST API Reference* throughout this guide and links to specific REST APIs to the API authentication and authorization for Amazon MQ section.
January 19, 2018	Updated the information in the Amazon MQ for ActiveMQ resources section.
January 18, 2018	Clarified and corrected the information in the Quotas in Amazon MQ section.
January 17, 2018	Reinstated the recommendation to prefer virtual destinations over durable subscriptions, with an improved explanation.
January 11, 2018	• The *Amazon MQ Developer Guide* is available in Kindle format, in addition to HTML and PDF. • Clarified and corrected information in the API authentication and a uthorization for Amazon MQ and Step 2: create a user and get your AWS credentials sections.
January 3, 2018	Added `DescribeConfigurationRevision` to the API authentication and authorization for Amazon MQ section.
December 15, 2017	Removed the recommendation against durable subscriptions from the Best Practices sections section.
December 8, 2017	• Added the Enable inbound connections prerequisite to the Connecting a Java application to your broker and Working Java Example sections. • Added the following note throughout this guide: Currently, you can't delete a configuration.
December 7, 2017	• Improved the code in the AmazonMQExample.java. • Added the API authentication and authorization for Amazon MQ section.

Date	Documentation Update

December 5, 2017	• Clarified and corrected information in the Monitoring Amazon MQ using CloudWatch section:
	  • Improved the metric descriptions.
	  • Added the Amazon MQ for ActiveMQ metrics and Dimensions for ActiveMQ broker metrics sub-sections.
	• Added the "Introducing Amazon MQ" video to the What is Amazon MQ? section.

December 4, 2017	• Clarified the following information in the Data Storage section: Storage capacity *per broker* is 200 GB. • Added the Prerequisites to the Working Java Example section. (The `activemq-client.jar` and `activemq-pool.jar` packages are required for the example to work. For more information, see Connecting a Java application to your broker).
December 1, 2017	• Updated and improved the screenshots in all the tutorials. • Clarified the following explanation throughout this guide: Making changes to a configuration revision or an ActiveMQ user does *not* apply the changes immediately. To apply your changes, you must wait for the next maintenance window or reboot the broker. For more information, see Amazon MQ broker configuration lifecycle.

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.