



Red Hat AMQ 7.7

Introducing Red Hat AMQ 7

Overview of Features and Components

Red Hat AMQ 7.7 Introducing Red Hat AMQ 7

Overview of Features and Components

Legal Notice

Copyright © 2020 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document highlights features and components of Red Hat AMQ 7. It also demonstrates common use cases and design patterns supported in this release.

Table of Contents

CHAPTER 1. ABOUT RED HAT AMQ 7	3
1.1. KEY FEATURES	3
Messaging at internet scale	3
Top-tier security and performance	3
Broad platform and language support	3
Focused on standards	3
Centralized management	3
CHAPTER 2. COMPONENT OVERVIEW	4
2.1. AMQ BROKER	4
2.2. AMQ INTERCONNECT	4
2.3. AMQ ONLINE	4
2.4. AMQ CLIENTS	4
AMQP clients	4
JMS clients	5
Adapters and libraries	5
2.5. COMPONENT COMPATIBILITY	5
CHAPTER 3. COMMON DEPLOYMENT PATTERNS	7
3.1. CENTRAL BROKER	7
3.2. ROUTED MESSAGING	7
3.3. HIGHLY AVAILABLE BROKERS	8
3.4. ROUTER PAIR BEHIND A LOAD BALANCER	9
3.5. ROUTER PAIR IN A DMZ	10
3.6. ROUTER PAIRS IN DIFFERENT DATA CENTERS	10

CHAPTER 1. ABOUT RED HAT AMQ 7

Red Hat AMQ provides fast, lightweight, and secure messaging for Internet-scale applications. AMQ Broker supports multiple protocols and fast message persistence. AMQ Interconnect leverages the AMQP protocol to distribute and scale your messaging resources across the network. AMQ Clients provides a suite of messaging APIs for multiple languages and platforms.

Think of the AMQ components as tools inside a toolbox. They can be used together or separately to build and maintain your messaging application, and AMQP is the glue in the toolbox that binds them together. AMQ components share a common management console, so you can manage them from a single interface.

1.1. KEY FEATURES

AMQ enables developers to build messaging applications that are fast, reliable, and easy to administer.

Messaging at internet scale

AMQ contains the tools to build advanced, multi-datacenter messaging networks. It can connect clients, brokers, and stand-alone services in a seamless messaging fabric.

Top-tier security and performance

AMQ offers modern SSL/TLS encryption and extensible SASL authentication. AMQ delivers fast, high-volume messaging and class-leading JMS performance.

Broad platform and language support

AMQ works with multiple languages and operating systems, so your diverse application components can communicate. AMQ supports C++, Java, JavaScript, Python, Ruby, and .NET applications, as well as Linux, Windows, and JVM-based environments.

Focused on standards

AMQ implements the Java JMS 1.1 and 2.0 API specifications. Its components support the ISO-standard AMQP 1.0 and MQTT messaging protocols, as well as STOMP and WebSocket.

Centralized management

With AMQ, you can administer all AMQ components from a single management interface. You can use JMX or the REST interface to manage servers programmatically.

CHAPTER 2. COMPONENT OVERVIEW

Red Hat AMQ consists of AMQ Broker, AMQ Interconnect, AMQ Online, and AMQ Clients. They work together to enable network communication in distributed applications.

- [AMQ Broker](#)
- [AMQ Interconnect](#)
- [AMQ Online](#)
- [AMQ Clients](#)

2.1. AMQ BROKER

AMQ Broker is a full-featured, message-oriented middleware broker. It offers advanced addressing and queuing, fast message persistence, and high availability. AMQ Broker supports multiple protocols and operating environments, enabling you to use your existing assets. AMQ Broker supports integration with Red Hat JBoss Enterprise Application Platform.

For more information, see [Getting Started with AMQ Broker](#).

2.2. AMQ INTERCONNECT

AMQ Interconnect provides flexible routing of messages between AMQP-enabled endpoints, including clients, brokers, and standalone services. With a single connection into a network of AMQ Interconnect routers, a client can exchange messages with any other endpoint connected to the network.

AMQ Interconnect does not use master-slave clusters for high availability. It is typically deployed in topologies of multiple routers with redundant network paths, which it uses to provide reliable connectivity. AMQ Interconnect can distribute messaging workloads across the network and achieve new levels of scale with very low latency.

For more information, see [Using AMQ Interconnect](#).

2.3. AMQ ONLINE

AMQ Online is an OpenShift-based system for delivering messaging as a managed service. With AMQ Online, administrators can configure a cloud-native, multi-tenant messaging service either in the cloud or on premise. Developers can provision messaging using the Red Hat AMQ Console. Multiple development teams can provision the brokers and queues from the AMQ Online console, without requiring each team to install, configure, deploy, maintain, or patch any software.

For more information, see [Using AMQ Online on OpenShift](#).

2.4. AMQ CLIENTS

AMQ Clients is a suite of AMQP 1.0 and JMS clients, adapters, and libraries. It includes JMS 2.0 support and new, event-driven APIs to enable integration into existing applications.

For more information, see [AMQ Clients Overview](#).

AMQP clients

- [AMQ C++](#)
- [AMQ JavaScript](#)
- [AMQ JMS \(Java\)](#)
- [AMQ .NET](#)
- [AMQ Python](#)
- [AMQ Ruby](#)

JMS clients

- [AMQ JMS \(AMQP 1.0\)](#)
- [AMQ Core Protocol JMS](#)
- [AMQ OpenWire JMS](#)

Adapters and libraries

- [AMQ JMS Pool](#)
- [AMQ Spring Boot Starter](#)

2.5. COMPONENT COMPATIBILITY

The following table lists the supported languages, platforms, and protocols of AMQ components. Note that any components supporting the same protocol can interoperate, even if their languages and platforms differ. For instance, AMQ Python can communicate with AMQ JMS.

Table 2.1. AMQ component compatibility

Component	Languages	Platforms	Protocols
AMQ Broker	-	JVM	AMQP 1.0, MQTT, OpenWire, STOMP, Core Protocol
AMQ Interconnect	-	Linux	AMQP 1.0
AMQ Online	-	OpenShift	AMQP 1.0
AMQ C++	C++	Linux, Windows	AMQP 1.0
AMQ JavaScript	JavaScript	Node.js, browsers	AMQP 1.0
AMQ JMS	Java	JVM	AMQP 1.0
AMQ .NET	C#	.NET	AMQP 1.0
AMQ Python	Python	Linux	AMQP 1.0

Component	Languages	Platforms	Protocols
AMQ Ruby	Ruby	Linux	AMQP 1.0
AMQ Spring Boot Starter	Java	JVM	AMQP 1.0
AMQ Core Protocol JMS	Java	JVM	Core Protocol
AMQ OpenWire JMS	Java	JVM	OpenWire
AMQ JMS Pool	Java	JVM	-

For more information, see [Red Hat AMQ 7 Supported Configurations](#) .


CHAPTER 3. COMMON DEPLOYMENT PATTERNS

Red Hat AMQ 7 can be set up in a large variety of topologies. The following are some of the common deployment patterns you can implement using AMQ components.

3.1. CENTRAL BROKER

The central broker pattern is relatively easy to set up and maintain. It is also relatively robust. Routes are typically local, because the broker and its clients are always within one network hop of each other, no matter how many nodes are added. This pattern is also known as *hub and spoke*, with the central broker as the hub and the clients the spokes.

Figure 3.1. Central broker pattern

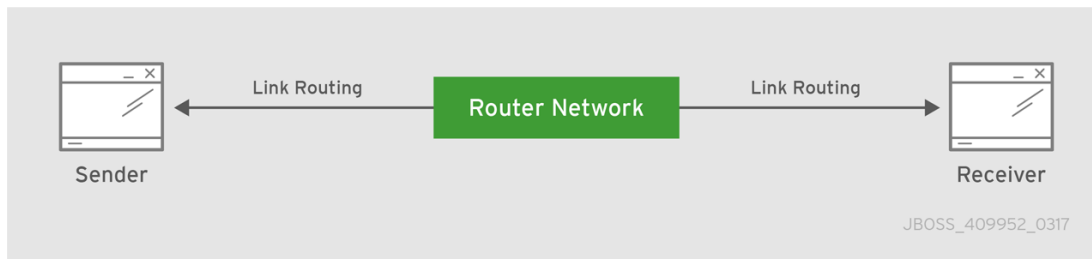


The only critical element is the central broker node. The focus of your maintenance efforts is on keeping this broker available to its clients.

3.2. ROUTED MESSAGING

When routing messages to remote destinations, the broker stores them in a local queue before forwarding them to their destination. However, sometimes an application requires sending request and response messages in real time, and having the broker store and forward messages is too costly. With AMQ you can use a router in place of a broker to avoid such costs. Unlike a broker, a router does not store messages before forwarding them to a destination. Instead, it works as a lightweight conduit and directly connects two endpoints.

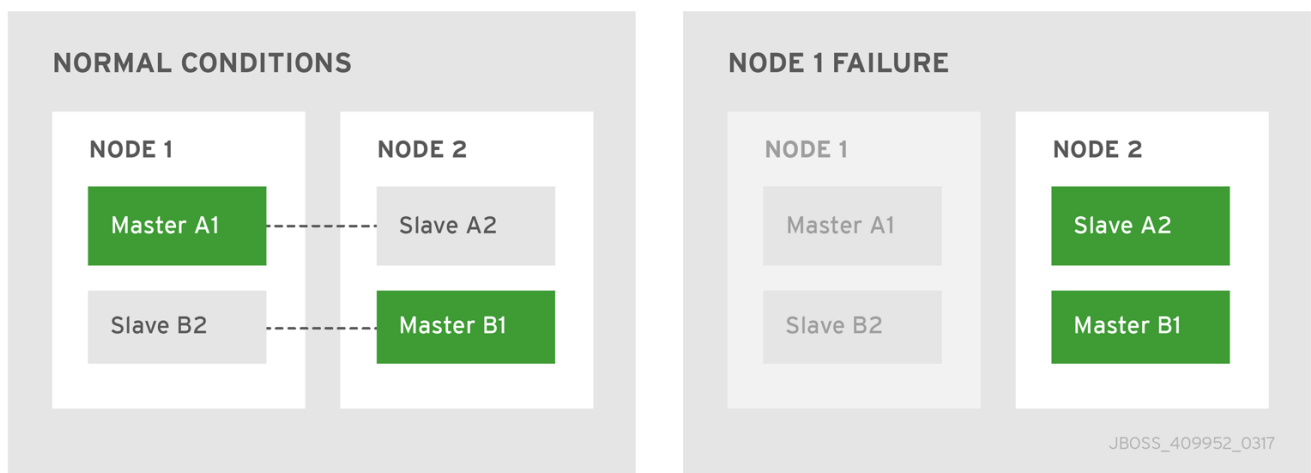
Figure 3.2. Brokerless routed messaging pattern



3.3. HIGHLY AVAILABLE BROKERS

To ensure brokers are available for their clients, deploy a highly available (HA) master-slave pair to create a backup group. You might, for example, deploy two master-slave groups on two nodes. Such a deployment would provide a backup for each active broker, as seen in the following diagram.

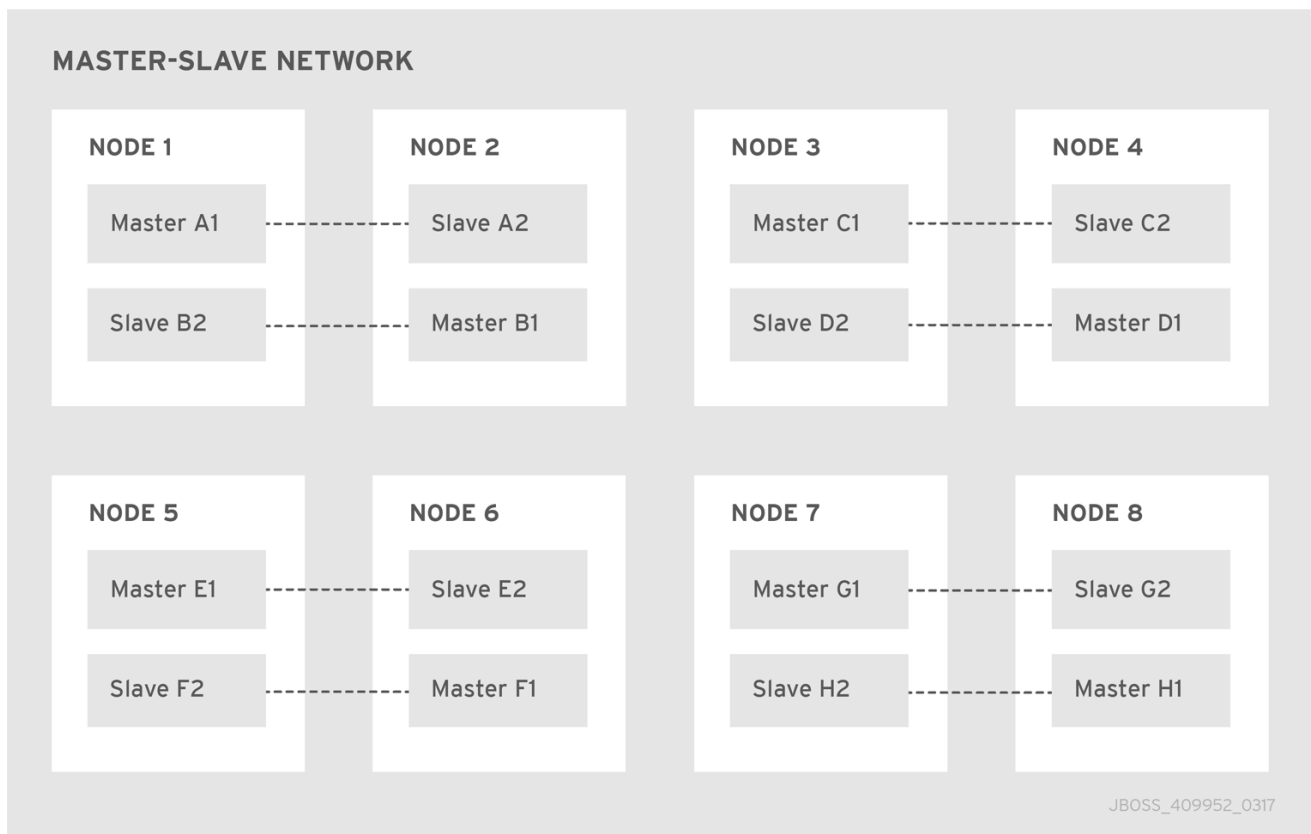
Figure 3.3. Master-slave pair



Under normal operating conditions one master broker is active on each node, which can be either a physical server or a virtual machine. If one node fails, the slave on the other node takes over. The result is two active brokers residing on the same healthy node.

By deploying master-slave pairs, you can scale out an entire network of such backup groups. Larger deployments of this type are useful for distributing the message processing load across many brokers. The broker network in the following diagram consists of eight master-slave groups distributed over eight nodes.

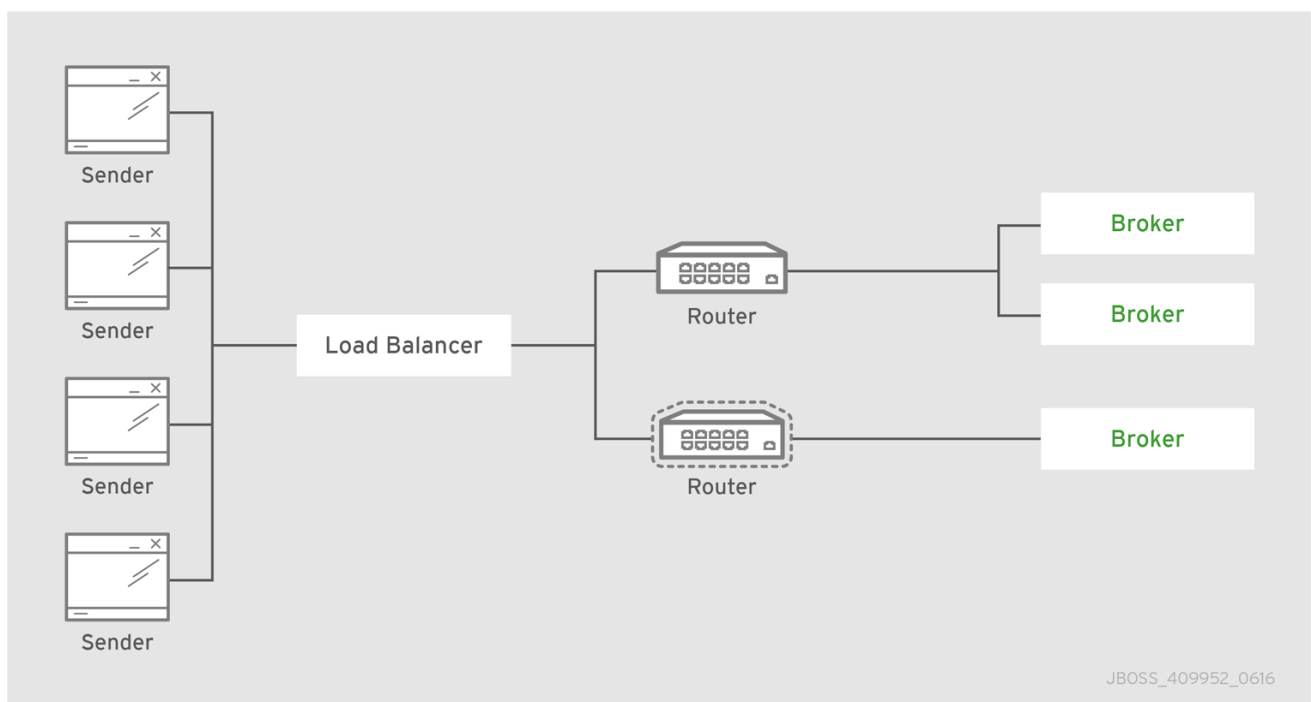
Figure 3.4. Master-slave network



3.4. ROUTER PAIR BEHIND A LOAD BALANCER

Deploying two routers behind a load balancer provides high availability, resiliency, and increased scalability for a single-datacenter deployment. Endpoints make their connections to a known URL, supported by the load balancer. Next, the load balancer spreads the incoming connections among the routers so that the connection and messaging load is distributed. If one of the routers fails, the endpoints connected to it will reconnect to the remaining active router.

Figure 3.5. Router pair behind a load balancer

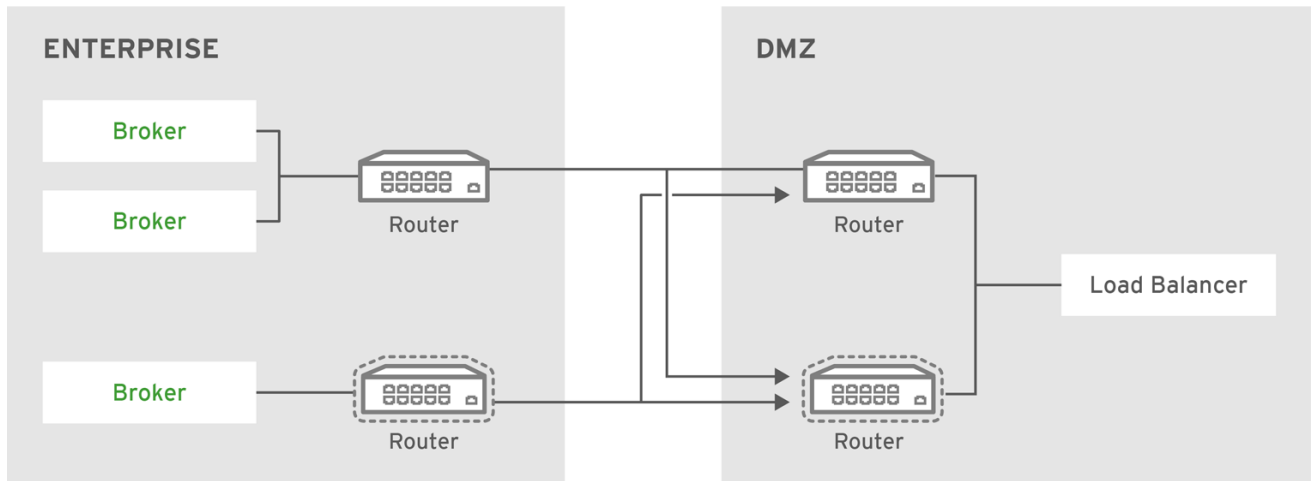


For even greater scalability, you can use a larger number of routers, three or four for example. Each router connects directly to all of the others.

3.5. ROUTER PAIR IN A DMZ

In this deployment architecture, the router network is providing a layer of protection and isolation between the clients in the outside world and the brokers backing an enterprise application.

Figure 3.6. Router pair in a DMZ



JBOSS_409952_0616


Important notes about the DMZ topology:

- Security for the connections within the deployment is separate from the security used for external clients. For example, your deployment might use a private Certificate Authority (CA) for internal security, issuing x.509 certificates to each router and broker for authentication, although external users might use a different, public CA.
- Inter-router connections between the enterprise and the DMZ are always established from the enterprise to the DMZ for security. Therefore, no connections are permitted from the outside into the enterprise. The AMQP protocol enables bi-directional communication after a connection is established, however.

3.6. ROUTER PAIRS IN DIFFERENT DATA CENTERS

You can use a more complex topology in a deployment of AMQ components that spans multiple locations. You can, for example, deploy a pair of load-balanced routers in each of four locations. You might include two backbone routers in the center to provide redundant connectivity between all locations. The following diagram is an example deployment spanning multiple locations.

Figure 3.7. Multiple interconnected routers



JBOSS_409952_0616

Revised on 2020-06-23 07:18:04 UTC