



# **Red Hat JBoss Enterprise Application Platform 7.0**

## **Configuring Messaging**

For Use with Red Hat JBoss Enterprise Application Platform 7.0



# Red Hat JBoss Enterprise Application Platform 7.0 Configuring Messaging

---

For Use with Red Hat JBoss Enterprise Application Platform 7.0

## Legal Notice

Copyright © 2018 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, 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 Software Collections 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 provides information for developers and administrators who want to develop and deploy messaging applications with JBoss EAP 7.0.

## Table of Contents

<b>PART I. ABOUT MESSAGING AND JBOSS EAP 7</b> .....	<b>7</b>
<b>CHAPTER 1. MESSAGING CONCEPTS</b> .....	<b>8</b>
1.1. MESSAGING SYSTEMS	8
1.2. MESSAGING STYLES	8
1.3. JAVA MESSAGING SERVICE (JMS)	8
1.4. JMS DESTINATIONS	9
<b>CHAPTER 2. THE INTEGRATED ACTIVEMQ ARTEMIS MESSAGING BROKER</b> .....	<b>10</b>
2.1. ACTIVEMQ ARTEMIS	10
2.2. APACHE ACTIVEMQ ARTEMIS CORE API AND JMS DESTINATIONS	10
<b>PART II. CONFIGURING SINGLE-NODE MESSAGING SYSTEMS</b> .....	<b>11</b>
<b>CHAPTER 3. GETTING STARTED</b> .....	<b>12</b>
3.1. USING THE HELLOWORLD-MDB QUICKSTART	12
Build and Deploy the helloworld-mdb Quickstart	12
3.2. OVERVIEW OF THE MESSAGING SUBSYSTEM CONFIGURATION	12
Connection Factories	13
Connectors and Acceptors	13
Socket Binding Groups	14
Messaging Security	14
Messaging Destinations	15
<b>CHAPTER 4. CONFIGURING MESSAGING DESTINATIONS</b> .....	<b>17</b>
4.1. ADDING A QUEUE	17
Reading a Queue's attributes	17
Attributes of a jms-queue	17
4.2. ADDING A TOPIC	18
Reading a Topic's attributes	18
Attributes of a jms-topic	19
4.3. JNDI ENTRIES AND CLIENTS	20
Management CLI Help	20
<b>CHAPTER 5. CONFIGURING LOGGING</b> .....	<b>21</b>
Configuring a Client for Logging	21
<b>CHAPTER 6. ADDRESS SETTINGS</b> .....	<b>23</b>
6.1. WILDCARD SYNTAX	23
6.2. DEFAULT ADDRESS-SETTING	23
Configuring Address Settings Using the Management CLI	24
Add a new address-setting	24
Edit an address-setting attribute	24
Read address-setting Attributes	24
Configuring Address Settings Using the Management Console	25
6.3. LAST-VALUE QUEUES	25
Configuring Last-value Queues	25
Using the Last-value Property	26
<b>CHAPTER 7. CONFIGURING SECURITY</b> .....	<b>27</b>
7.1. SECURING REMOTE CONNECTIONS	27
7.1.1. Securing the Transport	28
7.1.2. Securing a Remote Connector	28
7.2. SECURING DESTINATIONS	31

7.2.1. Role-Based Security for Addresses	32
Configuring Role-Based Security	32
7.2.1.1. Granting Unauthenticated Clients the guest Role	34
<b>CHAPTER 8. CONFIGURING THE MESSAGING TRANSPORTS</b>	<b>36</b>
8.1. ACCEPTOR AND CONNECTOR TYPES	36
8.2. ACCEPTORS	36
8.3. CONNECTORS	37
8.4. CONFIGURING ACCEPTORS AND CONNECTORS	38
8.5. CONNECTING TO A SERVER	40
8.5.1. JMS Connection Factories	40
8.5.2. Connecting to the Server Using JNDI	41
8.5.3. Connecting to the Server Using the Core API	42
ServerLocator	42
ClientSessionFactory	42
ClientSession	42
<b>CHAPTER 9. CONFIGURING CONNECTION FACTORIES</b>	<b>44</b>
Basic Connection Factories	44
Add a Connection Factory	44
Configure a Connection Factory	44
Remove a Connection Factory	44
Pooled Connection Factories	44
Add a Pooled Connection Factory	45
Configure a Pooled Connection Factory	45
Remove a Pooled Connection Factory	46
<b>CHAPTER 10. CONFIGURING PERSISTENCE</b>	<b>47</b>
10.1. ABOUT PERSISTENCE IN JBOSS EAP 7 MESSAGING	47
10.2. CONFIGURING THE BINDINGS AND JMS JOURNALS	49
10.3. CONFIGURING THE MESSAGE JOURNAL LOCATION	49
10.4. CONFIGURING MESSAGE JOURNAL ATTRIBUTES	50
10.5. NOTE ON DISABLING DISK WRITE CACHE	52
10.6. INSTALLING LIBAIO	53
10.7. CONFIGURING THE NFS SHARED STORE FOR MESSAGING	53
10.8. CONFIGURING JBOSS EAP MESSAGING FOR ZERO PERSISTENCE	54
10.9. IMPORTING AND EXPORTING JOURNAL DATA	54
<b>CHAPTER 11. CONFIGURING PAGING</b>	<b>55</b>
11.1. ABOUT PAGING	55
11.2. PAGE FILES	55
11.3. CONFIGURING THE PAGING DIRECTORY	55
11.4. CONFIGURING PAGING MODE	56
<b>CHAPTER 12. WORKING WITH LARGE MESSAGES</b>	<b>59</b>
12.1. STREAMING LARGE MESSAGES	59
Streaming Large Messages Using the Core API	59
Streaming Large Messages Over JMS	60
12.2. CONFIGURING LARGE MESSAGES	61
12.2.1. Configure Large Message Location	61
Configuring Large Message Size	62
Configuring Large Message Compression	62
12.2.2. Configuring Large Message Size Using the Core API	62
<b>CHAPTER 13. SCHEDULING MESSAGES</b>	<b>64</b>

<b>CHAPTER 14. TEMPORARY QUEUES AND RUNTIME QUEUES</b>	<b>65</b>
<b>CHAPTER 15. FILTER EXPRESSIONS AND MESSAGE SELECTORS</b>	<b>66</b>
<b>CHAPTER 16. CONFIGURING MESSAGE EXPIRY</b>	<b>68</b>
Set Message Expiry Using the Core API	68
Set Message Expiry Using JMS	68
16.1. EXPIRY ADDRESS	68
16.2. EXPIRY REAPER THREAD	68
<b>CHAPTER 17. CONFIGURING DELAYED REDELIVERY</b>	<b>70</b>
<b>CHAPTER 18. CONFIGURING DEAD LETTER ADDRESSES</b>	<b>71</b>
<b>CHAPTER 19. FLOW CONTROL</b>	<b>72</b>
19.1. CONSUMER FLOW CONTROL	72
Window-based flow control	72
Rate-limited flow control	73
19.2. PRODUCER FLOW CONTROL	74
Window-based flow control	74
Blocking producer window-based flow control	75
Rate-limited flow control	76
<b>CHAPTER 20. CONFIGURING PRE-ACKNOWLEDGMENTS</b>	<b>77</b>
20.1. CONFIGURING THE SERVER	77
20.2. CONFIGURING THE CLIENT	77
<b>CHAPTER 21. INTERCEPTORS</b>	<b>79</b>
21.1. IMPLEMENTING INTERCEPTORS	79
21.2. CONFIGURING INTERCEPTORS	79
<b>CHAPTER 22. MESSAGE GROUPING</b>	<b>81</b>
22.1. CONFIGURING MESSAGE GROUPS USING THE CORE API	81
22.2. CONFIGURING MESSAGE GROUPS USING JMS	81
<b>CHAPTER 23. DIVERTS</b>	<b>83</b>
23.1. EXCLUSIVE DIVERTS	84
23.2. NON-EXCLUSIVE DIVERTS	84
Creating diverts	84
<b>CHAPTER 24. THREAD MANAGEMENT</b>	<b>86</b>
24.1. SERVER SCHEDULED THREAD POOL	86
24.2. SERVER GENERAL PURPOSE THREAD POOL	86
24.3. EXPIRY REAPER THREAD	86
24.4. ASYNCHRONOUS IO	87
24.5. CLIENT THREAD MANAGEMENT	87
<b>CHAPTER 25. CONFIGURING DUPLICATE MESSAGE DETECTION</b>	<b>89</b>
25.1. USING DUPLICATE MESSAGE DETECTION FOR SENDING MESSAGES	89
25.2. CONFIGURING THE DUPLICATE ID CACHE	89
<b>CHAPTER 26. HANDLING SLOW CONSUMERS</b>	<b>91</b>
<b>PART III. CONFIGURING MULTI-NODE MESSAGING SYSTEMS</b>	<b>93</b>
<b>CHAPTER 27. CONFIGURING JMS BRIDGES</b>	<b>94</b>
27.1. QUALITY OF SERVICE	95

27.2. TIMEOUTS AND THE JMS BRIDGE	96
<b>CHAPTER 28. CONFIGURING CORE BRIDGES</b>	<b>97</b>
28.1. CONFIGURING A CORE BRIDGE FOR DUPLICATE DETECTION	97
<b>CHAPTER 29. CLUSTERS OVERVIEW</b>	<b>98</b>
29.1. SERVER DISCOVERY	98
29.1.1. Broadcast Groups	99
Configure a Broadcast Group Using UDP	99
Configure a Broadcast Group Using JGroups	100
Broadcast Group Attributes	100
29.1.2. Discovery Groups	100
29.1.2.1. Configure Discovery Groups on the Server	101
Configure a Discovery Group Using UDP	101
Configure a Discovery Group Using JGroups	101
Discovery Group Attributes	102
29.1.2.2. Configure Discovery Groups on the Client Side	102
Configure Client Discovery using JMS	102
Configure Client Discovery using the Core API	103
29.1.3. Static Discovery	104
Configuring a Cluster Connection	104
Configuring a Client Connection	104
Configuring Client Discovery Using JMS	104
Configuring Client Discovery Using the Core API	104
29.2. SERVER-SIDE MESSAGE LOAD BALANCING	105
Configuring the Cluster Connection	105
Configuring a Cluster Connection for Duplicate Detection	106
Cluster User Credentials	106
29.3. CLIENT-SIDE LOAD BALANCING	107
29.4. MESSAGE REDISTRIBUTION	108
29.5. CLUSTERED MESSAGE GROUPING	109
29.5.1. Best Practices for Clustered Message Grouping	110
<b>CHAPTER 30. HIGH AVAILABILITY</b>	<b>111</b>
30.1. LIVE / BACKUP PAIRS	111
30.2. HA POLICIES	111
30.3. DATA REPLICATION	112
30.3.1. Configuring Data Replication	114
30.3.2. All Replication Configuration	115
30.3.3. Preventing Cluster Connection Timeouts	116
30.3.4. Removing Old Journal Directories	117
Old Directories on the Live Server	118
30.3.5. Updating Dedicated Live and Backup Servers	118
30.3.6. Limitations of Data Replication: Split Brain Processing	118
30.4. SHARED STORE	119
30.4.1. Configuring a Shared Store	119
30.4.2. All Shared Store Configuration	120
30.5. FAILING BACK TO A LIVE SERVER	121
30.6. COLOCATED BACKUP SERVERS	122
30.6.1. Configuring Colocated Servers	122
30.7. FAILOVER MODES	127
30.7.1. Automatic Client Failover	127
Failing Over on the Initial Connection	129
About Server Replication	129



30.7.1.1. Handling Blocking Calls During Failover	129
30.7.1.2. Handling Failover With Transactions	130
30.7.1.3. Getting Notified of Connection Failure	131
30.7.2. Application-Level Failover	131
30.8. DETECTING DEAD CONNECTIONS	131
Cleaning up Dead Connection Resources on the Server	131
Closing Core Sessions or JMS Connections	133
Detecting Failure from the Client Side	133
Configuring Asynchronous Connection Execution	134
30.9. CLIENT RECONNECTION AND SESSION REATTACHMENT	134
Transparent Session Reattachment	134
Session Reconnection	135
Configuring Reconnection Attributes	135
ExceptionListeners and SessionFailureListeners	136
<b>CHAPTER 31. RESOURCE ADAPTERS</b>	<b>137</b>
31.1. ABOUT THE INTEGRATED ARTEMIS RESOURCE ADAPTER	137
Outbound Connection	137
Inbound Connections	137
31.2. USING THE INTEGRATED ARTEMIS RESOURCE ADAPTER FOR REMOTE CONNECTIONS	137
Configuring an MDB to use a pooled-connection-factory	138
Configuring the JMS destination	138
31.3. DEPLOYING AN A-MQ RESOURCE ADAPTER	139
31.3.1. Issues with the A-MQ 6 Resource Adapter	139
31.4. DEPLOYING THE IBM WEBSHERE® MQ RESOURCE ADAPTER	140
About IBM WebSphere® MQ	140
Summary	140
Prerequisites	140
Procedure to Deploy the IBM WebSphere® Resource Adapter	141
31.4.1. The Limitations and Known Problems for IBM WebSphere® MQ 7.5 Resource Adapter	143
31.5. DEPLOYING A GENERIC JMS RESOURCE ADAPTER	145
31.5.1. Configure a Generic JMS Resource Adapter for Use with a Third-party JMS Provider	146
<b>CHAPTER 32. BACKWARD AND FORWARD COMPATIBILITY</b>	<b>150</b>
32.1. FORWARD COMPATIBILITY	150
Management CLI migrate Operation	151
32.2. BACKWARD COMPATIBILITY	151
<b>PART IV. PERFORMANCE TUNING</b>	<b>153</b>
<b>CHAPTER 33. TUNING JMS</b>	<b>154</b>
<b>CHAPTER 34. TUNING PERSISTENCE</b>	<b>155</b>
<b>CHAPTER 35. OTHER TUNING OPTIONS</b>	<b>156</b>
<b>CHAPTER 36. AVOIDING ANTI-PATTERNS</b>	<b>157</b>
<b>APPENDIX A. REFERENCE MATERIAL</b>	<b>158</b>
A.1. ADDRESS SETTING ATTRIBUTES	158
A.2. CONNECTION FACTORY ATTRIBUTES	159
A.3. POOLED CONNECTION FACTORY ATTRIBUTES	161
A.4. CORE BRIDGE ATTRIBUTES	164
A.5. JMS BRIDGE ATTRIBUTES	165
A.6. CLUSTER CONNECTION ATTRIBUTES	167



## **PART I. ABOUT MESSAGING AND JBOSS EAP 7**

The messaging broker in JBoss EAP 6 was called HornetQ, a JBoss community project. The HornetQ codebase was donated to the Apache ActiveMQ project, and the HornetQ community joined that project to enhance the donated codebase and create a next-generation messaging broker. The result is Apache ActiveMQ Artemis, the messaging broker for JBoss EAP 7, providing messaging consolidation and backwards compatibility with JBoss EAP 6. While ActiveMQ Artemis retains protocol compatibility with the HornetQ broker in JBoss EAP 6, it also contains some smart new features. This guide will explore, and provide useful examples for, the many features of the ActiveMQ Artemis broker available in JBoss EAP 7.0.

# CHAPTER 1. MESSAGING CONCEPTS

## 1.1. MESSAGING SYSTEMS

Messaging systems allow you to loosely couple heterogeneous systems together with added reliability. Unlike systems based on a Remote Procedure Call (RPC) pattern, messaging systems primarily use an asynchronous message passing pattern with no tight relationship between requests and responses. Most messaging systems are flexible enough to also support a request-response mode if needed, but this is not a primary feature of messaging systems.

Messaging systems decouple a message's sender of messages from its consumer(s). In fact, the senders and consumers of messages are completely independent and know nothing of each other, which allows you to create flexible, loosely coupled systems. Large enterprises often use a messaging system to implement a message bus which loosely couples heterogeneous systems together. Message buses can form the core of an Enterprise Service Bus (ESB). Using a message bus to decouple disparate systems allows the system to grow and adapt more easily. It also allows more flexibility to add new systems or retire old ones since they don't have brittle dependencies on each other.

Messaging systems can also incorporate concepts such as delivery guarantees to ensure reliable messaging, transactions to aggregate the sending or consuming of multiple messages as a single unit of work, and durability to allow messages to survive server failure or restart.

## 1.2. MESSAGING STYLES

There are two kinds of messaging styles that most messaging systems support: the *point-to-point* pattern and the *publish-subscribe* pattern.

- **Point-to-Point Pattern**  
The point-to-point pattern involves sending a message to a single consumer listening on a queue. Once in the queue, the message is usually made persistent to guarantee delivery. Once the message has moved through the queue, the messaging system delivers it to a consumer. The consumer acknowledges the delivery of the message once it is processed. There can be multiple consumers listening on the same queue for the same message, but only one consumer will receive each message.
- **Publish-Subscribe Pattern**  
The publish-subscribe pattern allows senders to send messages to multiple consumers using a single destination. This destination is often known as a *topic*. Each topic can have multiple consumers, or subscribers, and unlike point-to-point messaging, every subscriber receives any message published to the topic.

Another interesting distinction is that subscribers can be durable. Durable subscriptions pass the server a unique identifier when connecting, which allows the server to identify and send any messages published to the topic since the last time the subscriber made a connection. Such messages are typically retained by the server even after a restart.

## 1.3. JAVA MESSAGING SERVICE (JMS)

The Java Messaging Service 2.0 (JMS) is defined in [JSR 343](#) and is a part of the Java EE 7 specification. JMS is a Java API that provides both point-to-point and publish-subscribe messaging styles. JMS also incorporates the use of transactions. JMS does not define a standard wire format so while vendors of JMS providers may all use the standard APIs, they may use different internal wire protocols to communicate between their clients and servers.

## 1.4. JMS DESTINATIONS

JMS destinations, along with JMS connection factories, are JMS administrative objects. Destinations are used by JMS clients for both producing and consuming messages. The destination allows the JMS client to specify the target when it produces messages and the source of messages when consuming messages. When using a publish-subscribe pattern, destinations are referred to as topics. When using a point-to-point pattern, destinations are referred to as queues.

Applications may use many different JMS destinations which are configured on the server side and usually accessed via JNDI.

## CHAPTER 2. THE INTEGRATED ACTIVEMQ ARTEMIS MESSAGING BROKER

### 2.1. ACTIVEMQ ARTEMIS

Apache ActiveMQ Artemis is an open source project for an asynchronous messaging system. It is high performance, embeddable, clustered and supports multiple protocols. JBoss EAP 7 uses Apache ActiveMQ Artemis as its JMS broker and is configured using the `messaging-activemq` subsystem. This fully replaces the HornetQ broker but retains protocol compatibility with JBoss EAP 6.

The core ActiveMQ Artemis is JMS-agnostic and provides a non-JMS API, which is referred to as the *core API*. ActiveMQ Artemis also provides a JMS client API which uses a facade layer to implement the JMS semantics on top of the core API. Essentially, JMS interactions are translated into core API operations on the client side using the JMS client API. From there, all operations are sent using the core client API and Apache ActiveMQ Artemis wire format. The server itself only uses the core API. For more details on the core API and its concepts, refer to the [ActiveMQ Artemis documentation](#).

### 2.2. APACHE ACTIVEMQ ARTEMIS CORE API AND JMS DESTINATIONS

Let's quickly discuss how JMS destinations are mapped to Apache ActiveMQ Artemis addresses.

Apache ActiveMQ Artemis core is JMS-agnostic. It does not have any concept of a JMS topic. A JMS topic is implemented in core as an address (the topic name) with zero or more queues bound to it. Each queue bound to that address represents a topic subscription. Likewise, a JMS queue is implemented as an address (the JMS queue name) with one single queue bound to it which represents the JMS queue.

By convention, all JMS queues map to core queues where the core queue name has the string "jms.queue." prepended to it. E.g. the JMS queue with the name "orders.europe" would map to the core queue with the name "jms.queue.orders.europe". The address at which the core queue is bound is also given by the core queue name.

For JMS topics the address at which the queues that represent the subscriptions are bound is given by prepending the string "jms.topic." to the name of the JMS topic. E.g. the JMS topic with name "news.europe" would map to the core address "jms.topic.news.europe"

In other words if you send a JMS message to a JMS queue with name "orders.europe" it will get routed on the server to any core queues bound to the address "jms.queue.orders.europe". If you send a JMS message to a JMS topic with name "news.europe" it will get routed on the server to any core queues bound to the address "jms.topic.news.europe".

If you want to configure settings for a JMS Queue with the name "orders.europe", you need to configure the corresponding core queue "jms.queue.orders.europe":

```
<!-- expired messages in JMS Queue "orders.europe" will be sent to the JMS
Queue "expiry.europe" -->
<address-setting match="jms.queue.orders.europe">
  <expiry-address>jms.queue.expiry.europe</expiry-address>
  ...
</address-setting>
```

## PART II. CONFIGURING SINGLE-NODE MESSAGING SYSTEMS

Part II begins with a guide to getting started with JBoss EAP 7 messaging by using the **helloworld-mdb** quickstart. Configuration options available to any installation follow, including topics such as security and persistence. For configuration relating to multiple installations of JBoss EAP 7, including topics such as clustering, high availability, and connecting to another server, see Part III, [Configuring Multi-Node Messaging Systems](#).

## CHAPTER 3. GETTING STARTED

### 3.1. USING THE HELLOWORLD-MDB QUICKSTART

The `helloworld-mdb` quickstart uses a simple message-driven bean to demonstrate basic Java EE 7 messaging features. Having the quickstart up and running as you [review the basic configuration](#) is an excellent way to introduce yourself to the features included with the JBoss EAP messaging server.

#### Build and Deploy the `helloworld-mdb` Quickstart

See the instructions in the `README.md` file provided with the quickstart for instructions on building and deploying the `helloworld-mdb` quickstart. You will need to start the JBoss EAP server specifying the `full` configuration, which contains the `messaging-activemq` subsystem. See the `README.md` file or the JBoss EAP [Configuration Guide](#) for details on starting JBoss EAP with a different configuration file.

### 3.2. OVERVIEW OF THE MESSAGING SUBSYSTEM CONFIGURATION

Default configuration for the `messaging-activemq` subsystem is included when starting the JBoss EAP server with the `full` or `full-ha` configuration. The `full-ha` option includes advanced configuration for features like [clustering](#) and [high availability](#).

Although not necessary, it is recommended that you [use the `helloworld-mdb` quickstart](#) as a working example to have running alongside this overview of the configuration.

For information on all settings available in the `messaging-activemq` subsystem, see the schema definitions located in the `EAP_HOME/docs/schema/` directory, or run the `read-resource-description` operation on the subsystem from the management CLI, as shown below.

```
/subsystem=messaging-activemq:read-resource-description(recursive=true)
```

The following extension in the server configuration file tells JBoss EAP to include the `messaging-activemq` subsystem as part of its runtime.

```
<extensions>
  ...
  <extension module="org.wildfly.extension.messaging-activemq"/>
  ...
</extensions>
```

The configuration for the `messaging-activemq` subsystem is contained within the `<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">` element.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    <cluster password="{jboss.messaging.cluster.password:CHANGE ME!!}" />
    <security-setting name="#">
      <role name="guest" send="true" consume="true" create-non-durable-queue="true" delete-non-durable-queue="true" />
    </security-setting>
    <address-setting name="#" dead-letter-address="jms.queue.DLQ" expiry-address="jms.queue.ExpiryQueue" max-size-bytes="10485760" page-size-bytes="2097152" message-counter-history-day-limit="10" redistribution-delay="1000" />
  </server>
</subsystem>
```



```

        <http-connector name="http-connector" socket-binding="http"
endpoint="http-acceptor"/>
        <http-connector name="http-connector-throughput" socket-
binding="http" endpoint="http-acceptor-throughput">
            <param name="batch-delay" value="50"/>
        </http-connector>
        <in-vm-connector name="in-vm" server-id="0"/>
        <http-acceptor name="http-acceptor" http-listener="default"/>
        <http-acceptor name="http-acceptor-throughput" http-
listener="default">
            <param name="batch-delay" value="50"/>
            <param name="direct-deliver" value="false"/>
        </http-acceptor>
        <in-vm-acceptor name="in-vm" server-id="0"/>
        <broadcast-group name="bg-group1" connectors="http-connector"
jgroups-channel="activemq-cluster"/>
        <discovery-group name="dg-group1" jgroups-channel="activemq-
cluster"/>
        <cluster-connection name="my-cluster" address="jms" connector-
name="http-connector" discovery-group="dg-group1"/>
        <jms-queue name="ExpiryQueue"
entries="java:/jms/queue/ExpiryQueue"/>
        <jms-queue name="DLQ" entries="java:/jms/queue/DLQ"/>
        <connection-factory name="InVmConnectionFactory" connectors="in-
vm" entries="java:/ConnectionFactory"/>
        <connection-factory name="RemoteConnectionFactory" ha="true"
block-on-acknowledge="true" reconnect-attempts="-1" connectors="http-
connector" entries="java:jboss/exported/jms/RemoteConnectionFactory"/>
        <pooled-connection-factory name="activemq-ra" transaction="xa"
connectors="in-vm" entries="java:/JmsXA
java:jboss/DefaultJMSConnectionFactory"/>
    </server>
</subsystem>

```

### Connection Factories

Messaging clients use a JMS **ConnectionFactory** object to make connections to the server. The default JBoss EAP configuration defines several connection factories. Note that there is a **<connection-factory>** for in-vm, http, and pooled connections.

```

    <connection-factory name="InVmConnectionFactory" connectors="in-vm"
entries="java:/ConnectionFactory"/>
    <connection-factory name="RemoteConnectionFactory" ha="true" block-on-
acknowledge="true" reconnect-attempts="-1" connectors="http-connector"
entries="java:jboss/exported/jms/RemoteConnectionFactory"/>
    <pooled-connection-factory name="activemq-ra" transaction="xa"
connectors="in-vm" entries="java:/JmsXA
java:jboss/DefaultJMSConnectionFactory"/>

```

See the [Configuring Connection Factories](#) section for more details.

### Connectors and Acceptors

Each JMS connection factory uses connectors to enable JMS-enabled communication from a client producer or consumer to a messaging server. The connector object defines the transport and parameters used to connect to the messaging server. Its counterpart is the acceptor object, which identifies the type of connections accepted by the messaging server.

The default JBoss EAP configuration defines several connectors and acceptors.

## Default Connectors

```
<http-connector name="http-connector" socket-binding="http"
endpoint="http-acceptor"/>
<http-connector name="http-connector-throughput" socket-binding="http"
endpoint="http-acceptor-throughput">
  <param name="batch-delay" value="50"/>
</http-connector>
<in-vm-connector name="in-vm" server-id="0"/>
```

## Default Acceptors

```
<http-acceptor name="http-acceptor" http-listener="default"/>
<http-acceptor name="http-acceptor-throughput" http-listener="default">
  <param name="batch-delay" value="50"/>
  <param name="direct-deliver" value="false"/>
</http-acceptor>
```

See the [Acceptors and Connectors](#) section for more details.

## Socket Binding Groups

The **socket-binding** attribute for the default connectors reference a socket binding named **http**. The http connector is used because JBoss EAP can multiplex inbound requests over standard web ports.

You can find this **socket-binding** as part of the **<socket-binding-group>** section elsewhere in the configuration file. Note how the configuration for the http and https socket bindings appear within the **<socket-binding-groups>** element:

```
<socket-binding-group name="standard-sockets" default-interface="public"
port-offset="{jboss.socket.binding.port-offset:0}">
  ...
  <socket-binding name="http" port="{jboss.http.port:8080}"/>
  <socket-binding name="https" port="{jboss.https.port:8443}"/>
  ...
</socket-binding-group>
```

For information on socket bindings, see [Configuring Socket Bindings](#) in the *JBoss EAP Configuration Guide*.

## Messaging Security

The **messaging-activemq** subsystem includes a single **security-setting** element when JBoss EAP is first installed:

```
<security-setting name="#">
  <role name="guest" delete-non-durable-queue="true" create-non-durable-
queue="true" consume="true" send="true"/>
</security-setting>
```

This declares that any user with the role **guest** can access any address on the server, as noted by the wildcard #. See [Configuring Address Settings](#) for more information on the wildcard syntax .

For more information on securing destinations and remote connections see [Configuring Messaging Security](#).

### Messaging Destinations

The **full** and **full-ha** configurations include two helpful queues that JBoss EAP can use to hold messages that have expired or that cannot be routed to their proper destination.

```
<jms-queue name="ExpiryQueue" entries="java:/jms/queue/ExpiryQueue"/>
<jms-queue name="DLQ" entries="java:/jms/queue/DLQ"/>
```

You can add your own messaging destinations in JBoss EAP using one of the following methods.

- Using the management CLI  
Use the following management CLI command to add a queue.

```
jms-queue add --queue-address=testQueue --
entries=queue/test,java:jboss/exported/jms/queue/test
```

Use the following management CLI command to add a topic.

```
jms-topic add --topic-address=testTopic --
entries=topic/test,java:jboss/exported/jms/topic/test
```

- Using the management console  
Messaging destinations can be configured from the management console by selecting the **Configuration** tab, navigating to the **Messaging - ActiveMQ** subsystem, and selecting **Queues/Topics** on the messaging provider.
- Defining your destinations using a Java EE 7 deployment descriptor or annotation.  
Starting with Java EE 7, deployment descriptors can include configuration for queues and topics. Below is a snippet from a Java EE 7 descriptor file that defines a JMS queue.

```
...
<jms-destination>
  <name>java:global/jms/MyQueue</name>
  <interfaceName>javax.jms.Queue</interfaceName>
  <destinationName>myQueue</destinationName>
</jms-destination>
...
```

For example, the message-driven beans in the **helloworld-mdb** quickstart contain annotations that define the queue and topic needed to run the application. Destinations created in this way will appear in the list of runtime queues. Use the management CLI to display the list of runtime queues. After deploying the quickstart the runtime queues it created will appear as below:

```
/subsystem=messaging-activemq/server=default/runtime-queue=*:read-
resource
{
  "outcome" => "success",
  "result" => [
    ...
    {
      "address" => [
```

```
        ("subsystem" => "messaging-activemq"),
        ("server" => "default"),
        ("runtime-queue" => "jms.queue.HelloWorldMDBQueue")
    ],
    "outcome" => "success",
    "result" => {"durable" => undefined}
},
...
{
    "address" => [
        ("subsystem" => "messaging-activemq"),
        ("server" => "default"),
        ("runtime-queue" => "jms.topic.HelloWorldMDBTopic")
    ],
    "outcome" => "success",
    "result" => {"durable" => undefined}
},
...
    ]
}
```

See [Configuring Messaging Destinations](#) for more detailed information.

## CHAPTER 4. CONFIGURING MESSAGING DESTINATIONS



### NOTE

Remember, configuring messaging destinations requires JBoss EAP to have messaging enabled. This functionality is enabled by default when running with the **standalone-full.xml** or **standalone-full-ha.xml** configuration files. The **domain.xml** configuration file also has messaging enabled.

### 4.1. ADDING A QUEUE

To add a JMS queue, use the **jms-queue** command from the management CLI:

```
jms-queue add --queue-address=myQueue --entries=[queue/myQueue
jms/queue/myQueue java:jboss/exported/jms/queue/myQueue]
```

Note how the **entries** attribute is a list containing multiple JNDI names separated by a single space. Also note the use of square brackets, **[ ]**, to enclose the list of JNDI names. The **queue-address** provides routing configuration, and **entries** provides a list of JNDI names that clients can use to look up the queue.

#### Reading a Queue's attributes

You can read a queue's configuration using the **jms-queue** command in the management CLI.

```
jms-queue read-resource --queue-address=myQueue
```

Alternatively, you can read a queue's configuration by accessing the **messaging-activemq** subsystem using the management CLI:

```
/subsystem=messaging-activemq/server=default/jms-queue=myQueue:read-
resource()
{
  "outcome" => "success",
  "result" => {
    "durable" => true,
    "entries" => ["queue/myQueue jms/queue/myQueue
java:jboss/exported/jms/queue/myQueue"],
    "legacy-entries" => undefined,
    "selector" => undefined
  }
}
```

#### Attributes of a **jms-queue**

The management CLI displays all the attributes of the **jms-queue** configuration element when given the following command:

```
/subsystem=messaging-activemq/server=default/jms-queue=*:read-resource-
description()
```

The table below provides all the attributes of a **jms-queue**:

Attribute	Description
consumer-count	The number of consumers consuming messages from this queue. Available at runtime.
dead-letter-address	The address to send dead messages to. See <a href="#">"Configuring Dead Letter Addresses"</a> for more information.
delivering-count	The number of messages that this queue is currently delivering to its consumers. Available at runtime.
durable	Whether the queue is durable or not. See <a href="#">"Messaging Styles"</a> for more information on durable subscriptions.
entries	The list of JNDI names the queue will be bound to. Required.
expiry-address	The address that will receive expired messages. See <a href="#">"Configuring Message Expiry"</a> for details.
legacy-entries	The JNDI names the queue will be bound to.
message-count	The number of messages currently in this queue. Available at runtime.
messages-added	The number of messages added to this queue since it was created. Available at runtime.
paused	Whether the queue is paused. Available at runtime.
queue-address	The queue address defines what address is used for routing messages. See <a href="#">"Configuring Address Settings"</a> for details on address settings. Required.
scheduled-count	The number of scheduled messages in this queue. Available at runtime.
selector	The queue selector. For more information on selectors see <a href="#">"Filter Expressions and Message Selectors"</a> .
temporary	Whether the queue is temporary. See <a href="#">"Temporary Queues and Runtime Queues"</a> for more information.

## 4.2. ADDING A TOPIC

Adding or reading a Topic is much like adding a Queue:

```
jms-topic add --topic-address=myTopic --entries=[topic/myTopic
jms/topic/myTopic java:jboss/exported/jms/topic/myTopic]
```

### Reading a Topic's attributes

Reading Topic attributes also has syntax similar to that used for a Queue:

■

```
jms-topic read-resource --topic-address=myTopic

entries
  topic/myTopic jms/topic/myTopic java:jboss/exported/jms/topic/myTopic
legacy-entries=n/a
```

```
/subsystem=messaging-activemq/server=default/jms-topic=myTopic:read-
resource
{
  "outcome" => "success",
  "result" => {
    "entries" => ["topic/myTopic jms/topic/myTopic
java:jboss/exported/jms/topic/myTopic"],
    "legacy-entries" => undefined
  }
}
```

### Attributes of a `jms-topic`

The management CLI displays all the attributes of the `jms-topic` configuration element when given the following command:

```
/subsystem=messaging-activemq/server=default/jms-topic=*:read-resource-
description()
```

The table below lists the attributes of a `jms-topic`:

Attribute	Description
delivering-count	The number of messages that this queue is currently delivering to its consumers. Available at runtime.
durable-message-count	The number of messages for all durable subscribers for this topic. Available at runtime.
durable-subscription-count	The number of durable subscribers for this topic. Available at runtime.
entries	The JNDI names the topic will be bound to. Required.
legacy-entries	The legacy JNDI names the topic will be bound to.
message-count	The number of messages currently in this queue. Available at runtime.
messages-added	The number of messages added to this queue since it was created. Available at runtime.
non-durable-message-count	The number of messages for all non-durable subscribers for this topic. Available at runtime.
non-durable-subscription-count	The number of non-durable subscribers for this topic. Available at runtime.

Attribute	Description
subscription-count	The number of (durable and non-durable) subscribers for this topic. Available at runtime.
temporary	Whether the topic is temporary.
topic-address	The address the topic points to. Required.

### 4.3. JNDI ENTRIES AND CLIENTS

A Queue or Topic must be bound to the `java:jboss/exported` namespace for a remote client to be able to look it up. The client must use the text after `java:jboss/exported/` when doing the lookup. For example, a queue named `testQueue` has for its `entries` the list `jms/queue/test` `java:jboss/exported/jms/queue/test`. A remote client wanting to send messages to `testQueue` would look up the queue using the string `jms/queue/test`. A local client on the other hand could look it up using `java:jboss/exported/jms/queue/test`, `java:jms/queue/test`, or more simply `jms/queue/test`.

#### Management CLI Help

You can find more information about the `jms-queue` and `jms-topic` commands by using the `--help` `--commands` flags:

```
jms-queue --help --commands
```

```
jms-topic --help --commands
```



## CHAPTER 5. CONFIGURING LOGGING

You can configure logging for the **messaging-activemq** subsystem by adding a log category in the JBoss EAP **logging** subsystem for **org.apache.activemq** and setting the desired log level. You can also configure a log handler for the category to configure how the log messages are recorded.

To see more information in the logs regarding XA transactions, change the log level of the **com.arjuna** category to a more verbose setting such as **TRACE** or **DEBUG**.

For more information on logging, including configuration for categories and other options, see the section on [logging](#) in the JBoss EAP *Configuration Guide*.

**Table 5.1. Logging Categories**

If you want logs for...	Use this category...
XA transactions	com.arjuna
All messaging activity	org.apache.activemq
Messaging Journal calls only	org.apache.activemq.artemis.journal
JMS calls only	org.apache.activemq.artemis.jms
Messaging utils calls only	org.apache.activemq.artemis.utils
Messaging core server only	org.apache.activemq.artemis.core.server

### Configuring a Client for Logging

Configure messaging clients by following these steps.

1. Download dependencies to the JBoss JMS client and log manager.  
If you are using Maven, add the following dependencies to your **pom.xml** file:

```
<dependencies>
  ...
  <dependency>
    <groupId>org.jboss.logmanager</groupId>
    <artifactId>jboss-logmanager</artifactId>
    <version>1.5.3.Final</version>
  </dependency>
  <dependency>
    <groupId>org.jboss.eap</groupId>
    <artifactId>wildfly-jms-client-bom</artifactId>
    <type>pom</type>
  </dependency>
  ...
</dependencies>
```

For more information, see the section on [using Maven with JBoss EAP](#) in the JBoss EAP *Development Guide*.

2. Create a properties file for the logger. Name it **logging.properties** and save it to a known location. Below is an example properties file. See the section on [logging](#) in the JBoss EAP *Development Guide* for more information on configuring logging options on the client side.

```
# Root logger option
loggers=org.jboss.logging,org.apache.activemq.artemis.core.server,org.apache.activemq.artemis.utils,org.apache.activemq.artemis.journal,org.apache.activemq.artemis.jms,org.apache.activemq.artemis.ra

# Root logger level
logger.level=INFO
# Apache ActiveMQ Artemis logger levels
logger.org.apache.activemq.artemis.jms.level=INFO
logger.org.apache.activemq.artemis.journal.level=INFO
logger.org.apache.activemq.artemis.utils.level=INFO
logger.org.apache.activemq.artemis.core.server.level=INFO

# Root logger handlers
logger.handlers=FILE

# File handler configuration
handler.FILE=org.jboss.logmanager.handlers.FileHandler
handler.FILE.level=FINE
handler.FILE.properties=autoFlush,fileName
handler.FILE.autoFlush=true
handler.FILE.fileName=activemq.log
handler.FILE.formatter=PATTERN

# Formatter pattern configuration
formatter.PATTERN=org.jboss.logmanager.formatters.PatternFormatter
formatter.PATTERN.properties=pattern
formatter.PATTERN.pattern=%d{HH:mm:ss,SSS} %-5p [%c] %s%E%n
```

3. Start the client with the expected parameters. When starting your client code using the **java** command, add the following parameters:

- a. Add the JBoss client and logger jars to the class path:

```
-cp /path/to/jboss-client.jar:/path/to/jboss-logmanager.jar
```

- b. Enable the JBoss logging manager:

```
-Djava.util.logging.manager=org.jboss.logmanager.LogManager
```

- c. Set the location of the logging properties file:

```
-Dlogging.configuration=/path/to/logging.properties
```

The full command to start the client will look something like the following example:

```
$ java -Djava.util.logging.manager=org.jboss.logmanager.LogManager -Dlogging.configuration=/path/to/logging.properties -cp /path/to/jboss-client.jar:/path/to/jboss-logmanager.jar org.example.MyClient
```

## CHAPTER 6. ADDRESS SETTINGS

The **messaging-activemq** subsystem has several configurable options which control aspects of how and when a message is delivered, how many attempts should be made, and when the message expires. These configuration options all exist within the **<address-setting>** configuration element. You can configure JBoss EAP to apply a single **<address-setting>** to multiple destinations by using a wildcard syntax.

### 6.1. WILDCARD SYNTAX

Wildcards can be used to match similar addresses with a single statement, much like how many systems use the asterisk character, **\***, to match multiple files or strings with a single query. The following table lists the special characters that can be used to define an **<address-setting>**.

**Table 6.1. JMS Wildcard Syntax**

Character	Description
. (a single period)	Denotes the space between words in a wildcard expression.
# (a pound or hash symbol)	Matches any sequence of zero or more words.
* (an asterisk)	Matches a single word.

The examples in the table below illustrate how wildcards are used to match a set of addresses.

**Table 6.2. JMS Wildcard Examples**

Example	Description
news.europe.#	Matches <b>news.europe</b> , <b>news.europe.sport</b> , <b>news.europe.politics.fr</b> , but not <b>news.usa</b> or <b>europe</b> .
news.*	Matches <b>news.europe</b> and <b>news.usa</b> , but not <b>news.europe.sport</b> .
news.*.sport	Matches <b>news.europe.sport</b> and <b>news.usa.sport</b> , but not <b>news.europe.fr.sport</b> .

### 6.2. DEFAULT ADDRESS-SETTING


Out of the box, JBoss EAP includes a single **address-setting** element as part of the configuration for the **messaging-activemq** subsystem:

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <address-setting
      name="#"
      dead-letter-address="jms.queue.DLQ"
```

```

    expiry-address="jms.queue.ExpiryQueue"
    max-size-bytes="10485760"
    page-size-bytes="2097152"
    message-counter-history-day-limit="10" />
    ...
</server>
</subsystem>

```



## NOTE

The use of a single # for the **name** attribute makes this default **address-setting** the configuration to be used for all destinations since # matches any address. You can continue to apply this catch-all configuration to all of your addresses, or you can add a new **address-setting** for each address or group of addresses that requires its own configuration set.

## Configuring Address Settings Using the Management CLI

Configuring address settings is done by using either the management CLI or the management console, but the management CLI exposes more of the configuration attributes for editing. See [Address Setting Attributes](#) in the appendix of this guide for the full list of attributes.

### Add a new address-setting

Use the **add** operation to create a new address setting if required. You can run this command from the root of the management CLI session, which in the following examples creates a new pattern named . You can include configuration attributes for the **address-setting**. Below, a new **address-setting** matching **news.europe.#** is created with its **dead-letter-address** attribute set to the queue **DLQ.news**, which was created beforehand. Examples for both a standalone server and a managed server domain using the **full** profile are shown respectively.

```
/subsystem=messaging-activemq/server=default/address-setting=news.europe.#/:add(dead-letter-address=DLQ.news)
```

```
/profile=full/subsystem=messaging-activemq/server=default/address-setting=news.europe.#/:add(dead-letter-address=DLQ.news)
```

### Edit an address-setting attribute

Use the **write-attribute** operation to write a new value to an attribute. You can use tab completion to help complete the command string as you type, as well as to expose the available attributes. The following example updates the **max-delivery-attempts** value to **10**.

```
/subsystem=messaging-activemq/server=default/address-setting=news.europe.#/:write-attribute(name=max-delivery-attempts,value=10)
```

```
/profile=full/subsystem=messaging-activemq/server=default/address-setting=news.europe.#/:write-attribute(name=max-delivery-attempts,value=10)
```

### Read address-setting Attributes

Confirm the values are changed by running the **read-resource** operation with the **include-runtime=true** parameter to expose all current values active in the server model.

```
/subsystem=messaging-activemq/server=default/address-
setting=news.europe.#/:read-resource(include-runtime=true)
```

```
/profile=full/subsystem=messaging-activemq/server=default/address-
setting=news.europe.#/:read-resource(include-runtime=true)
```

## Configuring Address Settings Using the Management Console

You can use the management console to create and review address settings by following these steps:


1. Log in to the management console.
2. Select the **Configuration** tab at the top of the screen. When running a managed domain select a profile from the **Profile** menu at the top left.
3. Expand the **Messaging - ActiveMQ** menu under **Subsystems**.
4. Select a messaging provider. In the default configuration, only one provider, called **default**, is shown.
5. Choose **Queues/Topics** from the drop down menu next to the selected messaging provider.
6. Click **Address Settings** on the menu appearing on the left of the console. A list of configured address settings appears on the right, as well as options to add, edit, and remove address settings. Add a new pattern by clicking **Add** or select an existing pattern and click **Edit** to update. Clicking **Remove** deletes the selected setting.

Remember that when adding a new pattern, for example **news.europe.#**, the **Pattern** field refers to the **name** attribute of the **address-setting** element. You enter this value when using the management CLI to read or write attributes.

You can edit only the **dead-letter-address**, **expiry-address**, **redelivery-delay**, and **max-delivery-attempts** attributes while using the management console. Other attributes must be configured using the management CLI.

## 6.3. LAST-VALUE QUEUES

Last-value queues are special queues which discard any messages when a newer message with the same value for a well-defined last-value property is put in the queue. In other words, a last-value queue only retains the last value. A typical application of a last-value queue might involve stock prices, where you are interested only in the latest price of a particular stock.



### IMPORTANT

Last-value queues will not work as expected if the queue has paging enabled. Be sure to [disable paging](#) before using a last-value queue.

### Configuring Last-value Queues

Last-value queues are defined within the **address-setting** configuration element:

```
<address-setting name="jms.queue.lastValueQueue" last-value-queue="true"
/>
```

Use the management CLI to read the value of **last-value-queue** for a given **address-setting**:

```

/subsystem=messaging-activemq/server=default/address-
setting=news.europe.#:read-attribute(name=last-value-queue)
{
    "outcome" => "success",
    "result" => false
}

```

The accepted values for **last-value-queue** are **true** or **false**. Use the management CLI to set either value, like so:

```

/subsystem=messaging-activemq/server=default/address-
setting=news.europe.#:write-attribute(name=last-value-queue,value=true)

```

```

/subsystem=messaging-activemq/server=default/address-
setting=news.asia.#:write-attribute(name=last-value-queue,value=false)

```

### Using the Last-value Property

The property name used to identify the last value is **\_AMQ\_LVQ\_NAME** (or the constant **Message.HDR\_LAST\_VALUE\_NAME** from the Core API). Let the following Java code illustrate how to use the last-value property.

- First, the publisher sends a message to the last-value queue

```

TextMessage message = session.createTextMessage("My 1st message with the
last-value property set");
message.setStringProperty("_AMQ_LVQ_NAME", "MY_MESSAGE");
producer.send(message);

```

- Then it sends another message to the queue using the same last-value

```

message = session.createTextMessage("My 2nd message with the last-value
property set");
message.setStringProperty("_AMQ_LVQ_NAME", "MY_MESSAGE");
producer.send(message);

```

- Next, the consumer receives the message with the last-value

```

TextMessage messageReceived = (TextMessage)messageConsumer.receive(5000);
System.out.format("Received message: %s\n", messageReceived.getText());

```

In the above example the client's output would be **"My 2nd message with the last-value property set"** since both messages set **\_AMQ\_LVQ\_NAME** to **"MY\_MESSAGE"**, and the second message was received in the queue after the first.

## CHAPTER 7. CONFIGURING SECURITY

### 7.1. SECURING REMOTE CONNECTIONS

JBoss EAP secures the **messaging-activemq** subsystem with the help of security realms and domains. See the JBoss EAP *Security Architecture* guide for more information on [security realms](#) and [security domains](#). The **messaging-activemq** subsystem is pre-configured to use the security realm named **ApplicationRealm** and the security domain named **other**.

The **ApplicationRealm** is defined near the top of the configuration file.

```
<management>
  <security-realms>
    ...
    <security-realm name="ApplicationRealm">
      <authentication>
        <local default-user="$local" allowed-users="*" skip-group-
loading="true"/>
        <properties
          path="application-users.properties"
          relative-to="jboss.server.config.dir" />
      </authentication>
      <authorization>
        <properties
          path="application-roles.properties"
          relative-to="jboss.server.config.dir" />
      </authorization>
    </security-realm>
  </security-realms>
  ...
</management>
```

As its name implies, **ApplicationRealm** is the default security realm for all application-focused subsystems in JBoss EAP such as the **messaging-activemq**, **undertow**, and **ejb3** subsystems. **ApplicationRealm** uses the local filesystem to store usernames and hashed passwords. For convenience JBoss EAP includes a script that you can use to add users to the **ApplicationRealm**. See [Default User Configuration](#) in the JBoss EAP *How To Configure Server Security* guide for details.

The **other** security domain is the default security domain for the application-related subsystems like **messaging-activemq**. It is not explicitly declared in the configuration; however, you can confirm which security domain is used by the **messaging-activemq** subsystem with the following management CLI command:

```
/subsystem=messaging-activemq/server=default:read-attribute(name=security-
domain)
{
  "outcome" => "success",
  "result" => "other"
}
```

You can also update which security domain is used:

```
/subsystem=messaging-activemq/server=default:write-attribute(name=security-domain, value=mySecurityDomain)
```

The JBoss EAP [How To Configure Server Security](#) guide has more information on how to create new security realms and domains. For now, it is worth noting how the **other** domain appears in the configuration:

```
<subsystem xmlns="urn:jboss:domain:security:1.2">
  <security-domains>
    <security-domain name="other" cache-type="default">
      <authentication>
        <login-module code="Remoting" flag="optional">
          <module-option name="password-stacking"
value="useFirstPass"/>
        </login-module>
        <login-module code="RealmDirect" flag="required">
          <module-option name="password-stacking"
value="useFirstPass"/>
        </login-module>
      </authentication>
    </security-domain>
    ...
  </security-domains>
</subsystem>
```

The 'other' domain uses two login-modules as its means of authentication. The first module, **Remoting**, authenticates remote EJB invocations, while the **RealmDirect** module uses the information store defined in a given realm to authenticate users. In this case the default realm **ApplicationRealm** is used, since no realm is declared. Each module has its **password-stacking** option set to **useFirstPass**, which tells the login-module to store the principal name and password of the authenticated user. See the JBoss EAP [Login Module Reference](#) for more details on the login modules and their options.

Role-based access is configured at the address level, see [Role Based Security for Addresses](#).

### 7.1.1. Securing the Transport

Using the default **http-connector** that comes bundled with JBoss EAP messaging and discussed in [Configuring the Messaging Transports](#), JBoss EAP is a great convenience in terms of port management. However it is not secured without further configuration. Fortunately securing the messaging transport is as easy as enabling normal web traffic for SSL/TLS. The steps to do this are provided in [Setting Up SSL/TLS for Applications](#) in JBoss EAP [How To Configure Server Security](#).

### 7.1.2. Securing a Remote Connector

If you are not using the default **http-connector** and have instead created your own **remote-connector** and **remote-acceptor** for TCP communications, you can configure each for SSL/TLS by using the properties in the table below. The properties appear in the configuration as part of the child **<param>** elements of the acceptor or connector.

Typically, a server owns its private SSL/TLS key and shares its public key with clients. In this scenario, the server defines the **key-store-path** and **key-store-password** parameters in a **remote-acceptor**. Since each client can have its trust store located at a different location, and be encrypted by a different password, specifying the **tust-store-path** and **trust-store-password** properties on



the **remote-connector** is not recommended. Instead, configure these parameters on the client side using the system properties **javax.net.ssl.trustStore** and **javax.net.ssl.trustStorePassword**. The only parameter which you need to configure for a **remote-connector** is **ssl-enabled=true**. However, if the server uses **remote-connector** to connect to another server, it makes sense in this case to set the **trust-store-path** and **trust-store-password** parameters of the **remote-connector**.

In the above use case, the **remote-acceptor** would be created using the following management CLI command:

```
/subsystem=messaging-activemq/server=default/remote-acceptor=mySslAcceptor:add(socket-binding=netty,params={ssl-enabled=true,key-store-path=path/to/server.jks,key-store-password=${VAULT::server-key::key-store-password::sharedKey}})
```

To create the **remote-connector** from the above use case, use the following management CLI command:

```
/subsystem=messaging-activemq/server=default/remote-connector=mySslConnector:add(socket-binding=netty,params={ssl-enabled=true})
```

The management CLI also allows you to add a parameter to an already existing **remote-acceptor** or **remote-connector** as well:

```
/subsystem=messaging-activemq/server=default/remote-connector=myOtherSslConnector:map-put(name=params,key=ssl-enabled,value=true)
```

Note that the **remote-acceptor** and **remote-connector** both reference a **socket-binding** to declare the port to be used for communication. See the [Overview of the Messaging Subsystem Configuration](#) for more information on socket bindings and their relationship to acceptors and connectors.

**Table 7.1. SSL/TLS-related Configuration Properties for the NettyConnectorFactory**

Property	Description
enabled-cipher-suites	Can be used to configure an acceptor or connector. This is a comma separated list of cipher suites used for SSL/TLS communication. The default value is null which means the JVM's default will be used.
enabled-protocols	Can be used to configure an acceptor or connector. This is a comma separated list of protocols used for SSL/TLS communication. The default value is null which means the JVM's default will be used.

Property	Description
key-store-password	<p>When used on an acceptor, this is the password for the server-side keystore.</p> <p>When used on a connector, this is the password for the client-side keystore. This is only relevant for a connector if you are using two-way SSL/TLS. Although this value can be configured on the server, it is downloaded and used by the client.</p> <p>If the client needs to use a different password from that set on the server, it can override the server-side setting by either using the standard <code>javax.net.ssl.keyStorePassword</code> system property. Use the <code>org.apache.activemq.ssl.keyStorePassword</code> property if another component on the client is already making use of the standard system property.</p>
key-store-path	<p>When used on an acceptor, this is the path to the SSL/TLS key store on the server which holds the server's certificates. Use for certificates either self-signed or signed by an authority.</p> <p>When used on a connector, this is the path to the client-side SSL/TLS key store which holds the client certificates. This is only relevant for a connector if you are using two-way SSL/TLS.</p> <p>Although this value is configured on the server, it is downloaded and used by the client. If the client needs to use a different path from that set on the server, it can override the server-side setting by using the standard <code>javax.net.ssl.keyStore</code> system property. Use the <code>org.apache.activemq.ssl.keyStore</code> system property if another component on the client is already making use of the standard property.</p>
key-store-provider	<p>Defines the format of the file in which keys are stored, PKCS11 or PKCS12 for example. The accepted values are JDK specific.</p>
needs-client-auth	<p>This property is only for an acceptor. It tells a client connecting to this acceptor that two-way SSL/TLS is required. Valid values are <b>true</b> or <b>false</b>. Default is <b>false</b>.</p>
ssl-enabled	<p>Must be <b>true</b> to enable SSL/TLS. Default is <b>false</b>.</p>

Property	Description
trust-store-password	<p>When used on an acceptor, this is the password for the server-side trust store. This is only relevant for an acceptor if you are using two-way SSL/TLS.</p> <p>When used on a connector, this is the password for the client-side truststore. Although this value can be configured on the server, it is downloaded and used by the client.</p> <p>If the client needs to use a different password from that set on the server, it can override the server-side setting by using either the standard <b>javax.net.ssl.trustStorePassword</b> system property. Use the <b>org.apache.activemq.ssl.trustStorePassword</b> system property if another component on the client is already making use of the standard property.</p>
trust-store-path	<p>When used on an acceptor, this is the path to the server-side SSL/TLS key store that holds the keys of all the clients that the server trusts. This is only relevant for an acceptor if you are using two-way SSL/TLS.</p> <p>When used on a connector, this is the path to the client-side SSL/TLS key store which holds the public keys of all the servers that the client trusts. Although this value can be configured on the server, it is downloaded and used by the client.</p> <p>If the client needs to use a different path from that set on the server, it can override the server-side setting by using either the standard <b>javax.net.ssl.trustStore</b> system property. Use the <b>org.apache.activemq.ssl.trustStore</b> system property if another component on the client is already making use of the standard system property.</p>
trust-store-provider	<p>Defines the format of the file in which keys are stored, PKCS11 or PKCS12 for example. The accepted values are JDK specific.</p>

## 7.2. SECURING DESTINATIONS

In addition to securing remote connections into the messaging server, you can also configure security around specific destinations. This is done by adding a security constraint using the **security-setting** configuration element. JBoss EAP messaging comes with a **security-setting** configured by default, as shown in the output from the following management CLI command:

```
/subsystem=messaging-activemq/server=default:read-resource(recursive=true)
{
  "outcome" => "success",
  "result" => {
    .....
    "security-setting" => {"#" => {"role" => {"guest" => {
      "consume" => true,
      "create-durable-queue" => false,
      "create-non-durable-queue" => true,
      "delete-durable-queue" => false,
```

```

        "delete-non-durable-queue" => true,
        "manage" => false,
        "send" => true
    }}}}
}

```

The **security-setting** option makes use of wildcards in the **name** field to handle which destinations to apply the security constraint. The value of a single # will match any address. For more information on using wildcards in security constraints, see [Role Based Security for Addresses](#).

### 7.2.1. Role-Based Security for Addresses

JBoss EAP messaging contains a flexible role-based security model for applying security to queues, based on their addresses.

The core JBoss EAP messaging server consists mainly of sets of queues bound to addresses. When a message is sent to an address, the server first looks up the set of queues that are bound to that address and then routes the message to the bound queues.

JBoss EAP messaging has a set of permissions that can be applied against queues based on their address. An exact string match on the address can be used or a wildcard match can be used using the wildcard characters # and \*. See [Address Settings](#) for more information on how to use the wildcard syntax.

You can create multiple roles for each **security-setting**, and there are 7 permission settings that can be applied to a role. Below is the complete list of the permissions available:

- **create-durable-queue** allows the role to create a durable queue under matching addresses.
- **delete-durable-queue** allows the role to delete a durable queue under matching addresses.
- **create-non-durable-queue** allows the role to create a non-durable queue under matching addresses.
- **delete-non-durable-queue** allows the role to delete a non-durable queue under matching addresses.
- **send** allows the role to send a message to matching addresses.
- **consume** allows the role to consume a message from a queue bound to matching addresses.
- **manage** allows the role to invoke management operations by sending management messages to the management address.

#### Configuring Role-Based Security

To start using role-based security for a **security-setting**, you first must create one. As an example, a **security-setting** of **news.europe.#** is created below. It would apply to any destination starting with **news.europe.**, such as **news.europe.fr** or **news.europe.tech.uk**.

```

/subsystem=messaging-activemq/server=default/security-
setting=news.europe.#:add()
{"outcome" => "success"}

```

Next, you add a role to the **security-setting** you created and declare permissions for it. In the example below, the **dev** role is created and given permissions to consume from, and send to, queues, as well as to create and delete non-durable queues. Because the default is **false**, you have to tell JBoss EAP only about the permissions you want to switch on.

```
/subsystem=messaging-activemq/server=default/security-
setting=news.europe.#/role=dev:add(consume=true,delete-non-durable-
queue=true,create-non-durable-queue=true,send=true)
{"outcome" => "success"}
```

To further illustrate the use of permissions, the example below creates an **admin** role and allows it to send management messages by switching on the **manage** permission. The permissions for creating and deleting durable queues are switched on as well:

```
/subsystem=messaging-activemq/server=default/security-
setting=news.europe.#/role=admin:add(manage=true,create-durable-
queue=true,delete-durable-queue=true)
{"outcome" => "success"}
```

To confirm the configuration of a **security-setting**, use the management CLI. Remember to use the **recursive=true** option to get the full display of permissions:

```
/subsystem=messaging-activemq/server=default:read-children-
resources(child-type=security-setting,recursive=true)
{
  "outcome" => "success",
  "result" => {
    "#" => {"role" => {"guest" => {
      "consume" => true,
      "create-durable-queue" => false,
      "create-non-durable-queue" => true,
      "delete-durable-queue" => false,
      "delete-non-durable-queue" => true,
      "manage" => false,
      "send" => true
    }}}},
    "news.europe.#" => {"role" => {
      "dev" => {
        "consume" => true,
        "create-durable-queue" => false,
        "create-non-durable-queue" => true,
        "delete-durable-queue" => false,
        "delete-non-durable-queue" => true,
        "manage" => false,
        "send" => true
      },
      "admin" => {
        "consume" => false,
        "create-durable-queue" => true,
        "create-non-durable-queue" => false,
        "delete-durable-queue" => true,
        "delete-non-durable-queue" => false,
        "manage" => true,
        "send" => false
      }
    }
  }
}
```

```

    }
  }}
}

```

Above, the permissions for addresses that start with string `news.europe.` are displayed in full by the management CLI. To summarize, only users who have the `admin` role can create or delete durable queues, while only users with the `dev` role can create or delete non-durable queues. Furthermore, users with the `dev` role can send or consume messages, but `admin` users cannot. They can, however, send management messages since their `manage` permission is set to `true`.

In cases where more than one match applies to a set of addresses the more specific match takes precedence. For example, the address `news.europe.tech.uk.#` is more specific than `news.europe.tech.#`. Because permissions are not inherited, you can effectively deny permissions in more specific `security-setting` blocks by simply not specifying them. Otherwise it would not be possible to deny permissions in sub-groups of addresses.

The mapping between a user and what roles they have is handled by the security manager. JBoss EAP ships with a user manager that reads user credentials from a file on disk, and can also plug into JAAS or JBoss EAP security.

For more information on configuring the security manager, please see the JBoss EAP [Security Architecture](#) guide.

### 7.2.1.1. Granting Unauthenticated Clients the guest Role

If you want JBoss EAP to automatically grant unauthenticated clients the `guest` role make the following two changes:

1. Add a new `module-option` to the `other` security domain. The new option, `unauthenticatedIdentity`, will tell JBoss EAP to grant `guest` access to unauthenticated clients. The recommended way to do this is by using the management CLI:

```

/subsystem=security/security-
domain=other/authentication=classic/login-module=RealmDirect:map-
put(name=module-options,key=unauthenticatedIdentity,value=guest)
{
  "outcome" => "success",
  "response-headers" => {
    "operation-requires-reload" => true,
    "process-state" => "reload-required"
  }
}

```

Note that the server requires a reload after issuing the command. You can confirm the new option by using the following management CLI command:

```

/subsystem=security/security-
domain=other/authentication=classic/login-module=RealmDirect:read-
resource()
{
  "outcome" => "success",
  "result" => {
    "code" => "RealmDirect",
    "flag" => "required",
    "module" => undefined,

```

```

        "module-options" => {
            "password-stacking" => "useFirstPass",
            "unauthenticatedIdentity" => "guest"
        }
    }
}

```

Also, your server configuration file should look something like this after the command executes:

```

<subsystem xmlns="urn:jboss:domain:security:1.2">
  <security-domains>
    <security-domain name="other" cache-type="default">
      <authentication>
        ...
        <login-module code="RealmDirect" flag="required">
          ...
          <module-option name="unauthenticatedIdentity"
value="guest"/>
          ...
        </login-module>
        ...
      </authentication>
    </security-domain>
    ...
  </security-domains>
</subsystem>

```

2. Uncomment the following line in the file **application-roles.properties** by deleting the # character. The file is located in **EAP\_HOME/standalone/configuration/** or **EAP\_HOME/domain/configuration/**, depending on whether you are using standalone servers or a domain controller respectively.

```
#guest=guest
```

Remote clients should now be able to access the server without needing to authenticate. They will be given the permissions associated with the **guest** role.

## CHAPTER 8. CONFIGURING THE MESSAGING TRANSPORTS

This section describes the concepts critical to understanding JBoss EAP messaging transports, specifically connectors and acceptors. Acceptors are used on the server to define how it can accept connections, while connectors are used by the client to define how it connects to a server. Each concept is discussed in turn and then a practical example shows how clients can make connections to a JBoss EAP messaging server, using JNDI or the Core API.

### 8.1. ACCEPTOR AND CONNECTOR TYPES

There are three main types of acceptor and connector defined in the configuration of JBoss EAP.

**in-vm:** In-vm is short for Intra Virtual Machine. Use this connector type when both the client and the server are running in the same JVM eg., Message Driven Beans (MDBs) running in the same instance of JBoss EAP.

**http:** Used when client and server are running in different JVMs. Uses the **undertow** subsystem's default port of **8080** and is thus able to multiplex messaging communications over HTTP. Red Hat recommends using the **http** connector when the client and server are running in different JVMs due to considerations such as port management, especially in a cloud environment.

**remote:** Remote transports are Netty-based components used for native TCP communication when the client and server are running in different JVMs. An alternative to **http** when it cannot be used.

A client must use a connector that is compatible with one of the server's acceptors. For example, only an **in-vm-connector** can connect to an **in-vm-acceptor**, and only a **http-connector** can connect to an **http-acceptor**, and so on.

You can have the management CLI list the attributes for a given acceptor or connector type using the **read-children-attributes** operation. For example, to see the attributes of all the **http-connectors** for the default messaging server you would enter:

```
/subsystem=messaging-activemq/server=default:read-children-
resources(child-type=http-connector,include-runtime=true)
```

The attributes of all the **http-acceptors** are read using a similar command:

```
/subsystem=messaging-activemq/server=default:read-children-
resources(child-type=http-acceptor,include-runtime=true)
```

The other acceptor and connector types follow the same syntax. Just provide **child-type** with the acceptor or connector type, for example, **remote-connector** or **in-vm-acceptor**.

### 8.2. ACCEPTORS

An acceptor defines which types of connection are accepted by the JBoss EAP integrated messaging server. You can define any number of acceptors per server. The sample configuration below is modified from the default **full-ha** configuration profile and provides an example of each acceptor type.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <http-acceptor name="http-acceptor" http-listener="default"/>
```



```

    <remote-acceptor name="legacy-messaging-acceptor" socket-
binding="legacy-messaging"/>
    <in-vm-acceptor name="in-vm" server-id="0"/>
    ...
</server>
</subsystem>

```

In the above configuration, the **http-acceptor** is using Undertow's default **http-listener** which listens on JBoss EAP's default http port, 8080. The **http-listener** is defined in the **undertow** subsystem:

```

<subsystem xmlns="urn:jboss:domain:undertow:3.0">
    ...
    <server name="default-server">
        <http-listener name="default" redirect-socket="https" socket-
binding="http"/>
        ...
    </server>
    ...
</subsystem>

```

Also note how the **remote-acceptor** above uses the **socket-binding** named **legacy-messaging**, which is defined later in the configuration as part of the server's default **socket-binding-group**.

```

<server xmlns="urn:jboss:domain:4.1">
    ...
    <socket-binding-group name="standard-sockets" default-interface="public"
port-offset="{jboss.socket.binding.port-offset:0}">
        ...
        <socket-binding name="legacy-messaging" port="5445"/>
        ...
    </socket-binding-group>
</server>

```

In this example, the **legacy-messaging socket-binding** binds JBoss EAP to port **5445**, and the **remote-acceptor** above claims the port on behalf of the **messaging-activemq** subsystem for use by legacy clients.

Lastly, the **in-vm-acceptor** uses a unique value for the **server-id** attribute so that this server instance can be distinguished from other servers that might be running in the same JVM.

## 8.3. CONNECTORS

A connector defines how to connect to an integrated JBoss EAP messaging server, and is used by a client to make connections.

You might wonder why connectors are defined on the server when they are actually used by the client. The reasons for this include:

- In some instances, the server might act as a client when it connects to another server. For example, one server might act as a bridge to another, or it might want to participate in a cluster. In such cases, the server needs to know how to connect to other servers, and that is defined by connectors.

- A server can provide connectors using a **ConnectionFactory** which is looked up by clients using JNDI, so creating connection to the server is simpler.

You can define any number of connectors per server. The sample configuration below is based on the **full-ha** configuration profile and includes connectors of each type.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <http-connector name="http-connector" endpoint="http-acceptor" socket-
binding="http"/>
    <remote-connector name="legacy-remoting-connector" socket-
binding="legacy-remoting"/>
    <in-vm-connector name="in-vm" server-id="0"/>
    ...
  </server>
</subsystem>
```

Like the **http-acceptor** from the **full-ha** profile, the **http-connector** uses the default **http-listener** defined by the **undertow** subsystem. The **endpoint** attribute declares which **http-acceptor** to connect to. In this case, the connector will connect to the default **http-acceptor**.

Also, note that the **remote-connector** references the same **socket-binding** as its **remote-acceptor** counterpart. Lastly, the **in-vm-connector** uses the same value for **server-id** as the **in-vm-acceptor** since they both run inside the same server instance.

## 8.4. CONFIGURING ACCEPTORS AND CONNECTORS

There are a number of configuration options for connectors and acceptors. They appear in the configuration as child **<param>** elements. Each **<param>** element includes a **name** and **value** attribute pair that is understood and used by the default Netty-based factory class responsible for instantiating a connector or acceptor.

In the management CLI, each remote connector or acceptor element includes an internal map of the parameter name and value pairs. For example, to add a new **param** to a **remote-connector** named **myRemote** use the following command:

```
/subsystem=messaging-activemq/server=default/remote-
connector=myRemote:map-put(name=params,key=foo,value=bar)
```

Retrieve parameter values using a similar syntax.

```
/subsystem=messaging-activemq/server=default/remote-
connector=myRemote:map-get(name=params,key=foo)
{
  "outcome" => "success",
  "result" => "bar"
}
```

You can also include parameters when you create an acceptor or connector, as in the example below.

```
/subsystem=messaging-activemq/server=default/remote-
connector=myRemote:add(socket-binding=mysocket,params={foo=bar,foo2=bar2})
```

Table 8.1. Transport Configuration Properties

Property	Description
batch-delay	Before writing packets to the transport, the messaging server can be configured to batch up writes for a maximum of <b>batch-delay</b> in milliseconds. This increases the overall throughput for very small messages by increasing average latency for message transfer. The default is 0.
direct-deliver	When a message arrives on the server and is delivered to waiting consumers, by default, the delivery is done on the same thread on which the message arrived. This gives good latency in environments with relatively small messages and a small number of consumers but reduces the throughput and latency. For highest throughput you can set this property as <b>false</b> . The default is <b>true</b> .
http-upgrade-enabled	Used by an <b>http-connector</b> to specify that it is using HTTP upgrade and therefore is multiplexing messaging traffic over HTTP. This property is set automatically by JBoss EAP to <b>true</b> when the <b>http-connector</b> is created and does not require an administrator.
http-upgrade-endpoint	Specifies the <b>http-acceptor</b> on the server-side to which the <b>http-connector</b> will connect. The connector will be multiplexed over HTTP and needs this info to find the relevant <b>http-acceptor</b> after the HTTP upgrade. This property is set automatically by JBoss EAP when the <b>http-connector</b> is created and does not require an administrator.
local-address	For a http or a remote connector, this is used to specify the local address which the client will use when connecting to the remote address. If a local address is not specified then the connector will use any available local address.
local-port	For a http or a remote connector, this is used to specify which local port the client will use when connecting to the remote address. If the local-port default is used (0) then the connector will let the system pick up an ephemeral port. Valid port values are <b>0</b> to <b>65535</b> .
nio-remoting-threads	If configured to use NIO, the messaging will by default use a number of threads equal to three times the number of cores (or hyper-threads) as reported by <b>Runtime.getRuntime().availableProcessors()</b> for processing incoming packets. To override this value, you can set a custom value for the number of threads. The default is <b>-1</b> .
tcp-no-delay	If this is <b>true</b> then Nagle's algorithm will be enabled. This algorithm helps improve the efficiency of TCP/IP networks by reducing the number of packets sent over a network. The default is <b>true</b> .
tcp-send-buffer-size	This parameter determines the size of the TCP send buffer in bytes. The default is <b>32768</b> .
tcp-receive-buffer-size	This parameter determines the size of the TCP receive buffer in bytes. The default is <b>32768</b> .

Property	Description
use-nio-global-worker-pool	This parameter will ensure all JMS connections share a single pool of Java threads, rather than each connection having its own pool. This serves to avoid exhausting the maximum number of processes on the operating system. The default is <b>true</b> .

## 8.5. CONNECTING TO A SERVER

If you want to connect a client to a server, you have to have a proper connector. There are two ways to do that. You could use a `ConnectionFactory` which is configured on the server and can be obtained via JNDI lookup. Alternatively, you could use the ActiveMQ Artemis core API and configure the whole `ConnectionFactory` on the client side.

### 8.5.1. JMS Connection Factories

Clients can use JNDI to look up `ConnectionFactory` objects which provide connections to the server. Connection Factories can expose each of the three types of connector:

A **connection-factory** referencing a **remote-connector** can be used by a remote client to send messages to or receive messages from the server (assuming the connection-factory has an appropriately exported entry). A **remote-connector** is associated with a **socket-binding** that tells the client using the **connection-factory** where to connect.

A **connection-factory** referencing an **in-vm-connector** is suitable to be used by a local client to either send messages to or receive messages from a local server. An **in-vm-connector** is associated with a **server-id** which tells the client using the **connection-factory** where to connect, since multiple messaging servers can run in a single JVM.

A **connection-factory** referencing a **http-connector** is suitable to be used by a remote client to send messages to or receive messages from the server by connecting to its HTTP port before upgrading to the messaging protocol. A **http-connector** is associated with the **socket-binding** that represents the HTTP socket, which by default is named **http**.

Since JMS 2.0, a default JMS connection factory is accessible to EE application under the JNDI name **java:comp/DefaultJMSConnectionFactory**. The **messaging-activemq** subsystem defines a **pooled-connection-factory** that is used to provide this default connection factory.

Below are the default connectors and connection factories that are included in the **full** configuration profile for JBoss EAP:

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    [...]
    <http-connector name="http-connector" socket-binding="http"
      endpoint="http-acceptor" />
    <http-connector name="http-connector-throughput" socket-binding="http"
      endpoint="http-acceptor-throughput">
      <param name="batch-delay" value="50"/>
    </http-connector>
    <in-vm-connector name="in-vm" server-id="0"/>
    [...]
    <connection-factory name="InVmConnectionFactory" connectors="in-vm"
```

```

entries="java:/ConnectionFactory" />
<pooled-connection-factory name="activemq-ra" transaction="xa"
connectors="in-vm" entries="java:/JmsXA
java:jboss/DefaultJMSConnectionFactory"/>
[...]
</server>
</subsystem>

```

The **entries** attribute of a factory specifies the JNDI names under which the factory will be exposed. Only JNDI names bound in the **java:jboss/exported** namespace are available to remote clients. If a **connection-factory** has an entry bound in the **java:jboss/exported** namespace a remote client would look-up the **connection-factory** using the text after **java:jboss/exported**. For example, the **RemoteConnectionFactory** is bound by default to **java:jboss/exported/jms/RemoteConnectionFactory** which means a remote client would look-up this connection-factory using **jms/RemoteConnectionFactory**. A **pooled-connection-factory** should not have any entry bound in the **java:jboss/exported** namespace because a **pooled-connection-factory** is not suitable for remote clients.

### 8.5.2. Connecting to the Server Using JNDI

If the client resides within the same JVM as the server, it can use the **in-vm** connector provided by the **InVmConnectionFactory**. Here is how the **InVmConnectionFactory** is typically configured, as found for example in **standalone-full.xml**.

```

<connection-factory
  name="InVmConnectionFactory"
  entries="java:/ConnectionFactory"
  connectors="in-vm" />

```

Note the value of the **entries** attribute. Clients using the **InVmConnectionFactory** should drop the leading **java:/** during lookup, as in the following example:

```

InitialContext ctx = new InitialContext();
ConnectionFactory cf = (ConnectionFactory)ctx.lookup("ConnectionFactory");
Connection connection = cf.createConnection();

```

Remote clients use the **RemoteConnectionFactory**, which is usually configured as below:

```

<connection-factory
  name="RemoteConnectionFactory"
  scheduled-thread-pool-max-size="10"
  entries="java:jboss/exported/jms/RemoteConnectionFactory"
  connectors="http-connector"/>

```

Remote clients should ignore the leading **java:jboss/exported/** of the value for **entries**, following the example of the code snippet below:

```

final Properties env = new Properties();
env.put(Context.INITIAL_CONTEXT_FACTORY,
"org.jboss.naming.remote.client.InitialContextFactory");
env.put(Context.PROVIDER_URL, "http-remoting://remotehost:8080");

```

```
InitialContext remotingCtx = new InitialContext(env);
ConnectionFactory cf = (ConnectionFactory)
remotingCtx.lookup("jms/RemoteConnectionFactory");
```

Note the value for the **PROVIDER\_URL** property and how the client is using the JBoss EAP http-remoting protocol. Note also how the client is using the **org.jboss.naming.remote.client.InitialContextFactory**, which implies the client has this class and its encompassing client jar somewhere in the classpath. For maven projects, this can be achieved by including the following dependency:

```
<dependencies>
  <dependency>
    <groupId>org.wildfly</groupId>
    <artifactId>wildfly-jms-client-bom</artifactId>
    <type>pom</type>
  </dependency>
</dependencies>
```

### 8.5.3. Connecting to the Server Using the Core API

You can use the Core API to make client connections without needing a JNDI lookup. Clients using the Core API require a client jar in their classpath, just as JNDI-based clients.

#### ServerLocator

Clients use **ServerLocator** instances to create **ClientSessionFactory** instances. As their name implies, **ServerLocator** instances are used to locate servers and create connections to them.

In JMS terms think of a **ServerLocator** in the same way you would a JMS Connection Factory.

**ServerLocator** instances are created using the **ActiveMQClient** factory class.

```
ServerLocator locator = ActiveMQClient.createServerLocatorWithoutHA(new
TransportConfiguration(InVMConnectorFactory.class.getName()));
```

#### ClientSessionFactory

Clients use a **ClientSessionFactory** to create **ClientSession** instances, which are basically connections to a server. In JMS terms think of them as JMS connections.

**ClientSessionFactory** instances are created using the **ServerLocator** class.

```
ClientSessionFactory factory = locator.createClientSessionFactory();
```

#### ClientSession

A client uses a **ClientSession** for consuming and producing messages and for grouping them in transactions. **ClientSession** instances can support both transactional and non transactional semantics and also provide an XAResource interface so messaging operations can be performed as part of a JTA transaction.

**ClientSession** instances group **ClientConsumers** and **ClientProducers**.

```
ClientSession session = factory.createSession();
```

The simple example below highlights some of what was just discussed:

```
ServerLocator locator = ActiveMQClient.createServerLocatorWithoutHA(
    new TransportConfiguration( InVMConnectorFactory.class.getName()));

// In this simple example, we just use one session for both
// producing and consuming
ClientSessionFactory factory = locator.createClientSessionFactory();
ClientSession session = factory.createSession();

// A producer is associated with an address ...
ClientProducer producer = session.createProducer("example");
ClientMessage message = session.createMessage(true);
message.getBodyBuffer().writeString("Hello");

// We need a queue attached to the address ...
session.createQueue("example", "example", true);

// And a consumer attached to the queue ...
ClientConsumer consumer = session.createConsumer("example");

// Once we have a queue, we can send the message ...
producer.send(message);

// We need to start the session before we can -receive- messages ...
session.start();
ClientMessage msgReceived = consumer.receive();

System.out.println("message = " +
    msgReceived.getBodyBuffer().readString());

session.close();
```

## CHAPTER 9. CONFIGURING CONNECTION FACTORIES

By default, the JBoss EAP `messaging-activemq` subsystem provides the `InVmConnectionFactory` and `RemoteConnectionFactory` [connection factories](#), as well as the `activemq-ra` [pooled connection factory](#).

### Basic Connection Factories

`InVmConnectionFactory` references an `in-vm-connector` and can be used to send and receive messages when both the client and server are running in the same JVM. `RemoteConnectionFactory` references an `http-connector` and can be used to send and receive messages over HTTP when the client and server are running in different JVMs.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <connection-factory name="InVmConnectionFactory" connectors="in-vm"
entries="java:/ConnectionFactory"/>
    <connection-factory name="RemoteConnectionFactory" connectors="http-
connector" entries="java:jboss/exported/jms/RemoteConnectionFactory"/>
    ...
  </server>
</subsystem>
```

For more information on the different types of connectors, see the [Acceptors and Connectors](#) section.

### Add a Connection Factory

You can add a new connection factory using the following management CLI command. When adding a connection factory, you must provide the `connectors` and the JNDI `entries`.

```
/subsystem=messaging-activemq/server=default/connection-
factory=MyConnectionFactory:add(entries=
[java:/MyConnectionFactory],connectors=[in-vm])
```

### Configure a Connection Factory

You can update a connection factory's settings using the management CLI.

```
/subsystem=messaging-activemq/server=default/connection-
factory=MyConnectionFactory:write-attribute(name=thread-pool-max-
size,value=40)
```

For information on the available attributes for a connection factory, see [Connection Factory Attributes](#).

### Remove a Connection Factory

You can remove a connection factory using the management CLI.

```
/subsystem=messaging-activemq/server=default/connection-
factory=MyConnectionFactory:remove
```

### Pooled Connection Factories

The JBoss EAP `messaging-activemq` subsystem provides a pooled connection factory that allows you to configure the inbound and outbound connectors of the integrated ActiveMQ Artemis resource adapter. For more information on configuring a `pooled-connection-factory` to connect to a remote ActiveMQ Artemis server, see [Using the Integrated Resource Adapter for Remote Connections](#).



```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <pooled-connection-factory name="activemq-ra" transaction="xa"
entries="java:/JmsXA java:jboss/DefaultJMSConnectionFactory"
connectors="in-vm"/>
  </server>
</subsystem>
```

There are several unique characteristics of a pooled connection factory:

- It is only available to local clients, though it can be configured to point to a remote server. For more information on connecting to a remote ActiveMQ Artemis server, see [Using the Integrated Artemis Resource Adapter for Remote Connections](#).
- It should only be used to send messages when looked up in JNDI or injected.
- It can be configured to use security credentials, which is useful if it is pointing to a secured remote server.
- Resources acquired from it will be automatically enlisted in any ongoing JTA transactions.

### Add a Pooled Connection Factory

You can add a new pooled connection factory using the following management CLI command. When adding a connection factory, you must provide the **connectors** and the JNDI **entries**.

```
/subsystem=messaging-activemq/server=default/pooled-connection-
factory=MyPooledConnectionFactory:add(entries=
[java:/MyPooledConnectionFactory],connectors=[in-vm])
```

### Configure a Pooled Connection Factory

You can update a pooled connection factory's settings using the management CLI.

```
/subsystem=messaging-activemq/server=default/pooled-connection-
factory=MyPooledConnectionFactory:write-attribute(name=max-retry-
interval,value=3000)
```

For information on the available attributes for a pooled connection factory, see [Pooled Connection Factory Attributes](#).

You can disable the recording of enlistment traces for this pooled connection factory using the management CLI by setting the **enlistment-trace** attribute to **false**.

```
/subsystem=messaging-activemq/server=default/pooled-connection-
factory=MyPooledConnectionFactory:write-attribute(name=enlistment-
trace,value=false)
```

**WARNING**

Disabling the enlistment trace will make tracking down errors during transaction enlistment more difficult.

You can also configure the managed connection pool implementation used by the pooled connection factory. For more information, see the [Configure Managed Connection Pools](#) section of the JBoss EAP *Configuration Guide*.

**Remove a Pooled Connection Factory**

You can remove a pooled connection factory using the management CLI.

```
/subsystem=messaging-activemq/server=default/pooled-connection-  
factory=MyPooledConnectionFactory:remove
```

## CHAPTER 10. CONFIGURING PERSISTENCE

### 10.1. ABOUT PERSISTENCE IN JBOSS EAP 7 MESSAGING

JBoss EAP messaging handles its own persistence. It ships with a high-performance journal, which is optimized for messaging.

The JBoss EAP messaging journal has a configurable file size and is append only, which improves performance by enabling single write operations. It consists of a set of files on disk, which are initially pre-created to a fixed size and filled with padding. As server operations (add message, delete message, update message, etc.) are performed, records of the operations are appended to the journal until the journal file is full, at which point the next journal file is used.

A sophisticated garbage collection algorithm determines whether journal files can be reclaimed and re-used when all of their data has been deleted. A compaction algorithm removes dead space from journal files and compresses the data.

The journal also fully supports both local and XA transactions.

The majority of the journal is written in Java, but interaction with the file system has been abstracted to allow different pluggable implementations. The two implementations shipped with JBoss EAP messaging are:

#### Java New I/O (NIO)

This implementation uses standard Java NIO to interface with the file system. It provides extremely good performance and runs on any platform with a Java 6 or later runtime. Note that JBoss EAP 7 requires Java 8. Using NIO is supported on any operating system that JBoss EAP supports.

#### Linux Asynchronous IO (ASYNCIO)

This implementation uses a native code wrapper to talk to the Linux asynchronous IO library (ASYNCIO). This implementation removes the need for explicit synchronization. ASYNCIO typically provides better performance than Java NIO.

The following file systems have been tested and are supported only on Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 when using the **libaio** natives. They are not tested and are not supported on other operating systems.

- EXT4
- XFS
- NFSv4
- GFS2

The following table lists the HA shared store file systems that have been tested, both with and without the **libaio** natives, and whether they are supported.

Operating System	File System	Supported Using <b>libaio</b> Natives? (journal-type="ASYNCIO")	Supported Without Using <b>libaio</b> Natives? (journal-type="NIO")
Red Hat Enterprise Linux 6	NFSv4	Yes	Yes

Operating System	File System	Supported Using <code>libaio</code> Natives? ( <code>journal-type="ASYNCIO"</code> )	Supported Without Using <code>libaio</code> Natives? ( <code>journal-type="NIO"</code> )
Red Hat Enterprise Linux 7	NFSv4	Yes	Yes
Red Hat Enterprise Linux 6	GFS2	Yes	No
Red Hat Enterprise Linux 7	GFS2	Yes	No

The standard JBoss EAP messaging core server uses the following journal instances:

### Bindings journal

This journal is used to store bindings related data, including the set of queues that are deployed on the server and their attributes. It also stores data such as id sequence counters.

The bindings journal is always a NIO journal as it is typically low throughput compared to the message journal.

The files on this journal are prefixed as `activemq-bindings`. Each file has a `bindings` extension. File size is 1048576, and it is located at the `bindings` folder.

### JMS journal

This journal instance stores all JMS related data, such as any JMS queues, topics, connection factories and any JNDI bindings for these resources.

Any JMS Resource created via the management API will be persisted to this journal. Any resource configured via configuration files will not. The JMS Journal will only be created if JMS is being used.

The files on this journal are prefixed as `activemq-jms`. Each file has a `jms` extension. File size is 1048576, and it is located at the `bindings` folder.

### Message journal

This journal instance stores all message related data, including the message themselves and also duplicate-id caches.

By default JBoss EAP messaging will try to use an ASYNCIO journal. If ASYNCIO is not available, for example the platform is not Linux with the correct kernel version or ASYNCIO has not been installed then it will automatically fall back to using Java NIO which is available on any Java platform.

The files on this journal are prefixed as `activemq-data`. Each file has an `amq` extension. File size is by default 10485760 (configurable), and it is located at the `journal` folder.

For large messages, JBoss EAP messaging persists them outside the message journal. This is discussed in the section on [Large Messages](#).

JBoss EAP messaging can also be configured to page messages to disk in low memory situations. This is discussed in the [Paging section](#).

If no persistence is required at all, JBoss EAP messaging can also be configured not to persist any data at all to storage as discussed in the [Configuring JBoss EAP Messaging for Zero Persistence](#) section.

## 10.2. CONFIGURING THE BINDINGS AND JMS JOURNALS

Because the bindings journal shares its configuration with the JMS journal, you can read the current configuration for both by using the single management CLI command below. The output is also included to highlight default configuration.

```
/subsystem=messaging-activemq/server=default/path=bindings-directory:read-resource
{
  "outcome" => "success",
  "result" => {
    "path" => "activemq/bindings",
    "relative-to" => "jboss.server.data.dir"
  }
}
```

Note that by default the **path** to the journal is **activemq/bindings**. You can change the location for **path** by using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/path=bindings-directory:write-attribute(name=path,value=<PATH_LOCATION>)
```

Also note the **relative-to** attribute in the output above. When **relative-to** is used, the value of the **path** attribute is treated as relative to the file path specified by **relative-to**. By default this value is the JBoss EAP property **jboss.server.data.dir**. For standalone servers, **jboss.server.data.dir** is located at **EAP\_HOME/standalone/data**. For domains, each server will have its own **serverX/data/activemq** directory located under **EAP\_HOME/domain/servers**. You can change the value of **relative-to** using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/path=bindings-directory:write-attribute(name=relative-to,value=<RELATIVE_LOCATION>)
```

By default, JBoss EAP is configured to automatically create the bindings directory if it does not exist. Use the following management CLI command to toggle this behavior.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=create-bindings-dir,value=<TRUE_FALSE>)
```

Setting **value** to **true** will enable automatic directory creation. Setting **value** to **false** will disable it.

## 10.3. CONFIGURING THE MESSAGE JOURNAL LOCATION

You can read the location information for the message journal by using the management CLI command below. The output is also included to highlight default configuration.

```
/subsystem=messaging-activemq/server=default/path=journal-directory:read-resource
{
  "outcome" => "success",
```

```
"result" => {
  "path" => "activemq/journal",
  "relative-to" => "jboss.server.data.dir"
}
}
```

Note that by default the **path** to the journal is **activemq/journal**. You can change the location for **path** by using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/path=journal-directory:write-attribute(name=path,value=<PATH_LOCATION>)
```



## NOTE

For the best performance, Red Hat recommends that the journal be located on its own physical volume in order to minimize disk head movement. If the journal is on a volume which is shared with other processes which might be writing other files, such as a bindings journal, database, or transaction coordinator, then the disk head may well be moving rapidly between these files as it writes them, thus drastically reducing performance.

Also note the **relative-to** attribute in the output above. When **relative-to** is used, the value of the **path** attribute is treated as relative to the file path specified by **relative-to**. By default this value is the JBoss EAP property **jboss.server.data.dir**. For standalone servers, **jboss.server.data.dir** is located at **EAP\_HOME/standalone/data**. For domains, each server will have its own **serverX/data/activemq** directory located under **EAP\_HOME/domain/servers**. You can change the value of **relative-to** using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/path=journal-directory:write-attribute(name=relative-to,value=<RELATIVE_LOCATION>)
```

By default, JBoss EAP is configured to automatically create the journal directory if it does not exist. Use the following management CLI command to toggle this behavior.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=create-journal-dir,value=<TRUE_FALSE>)
```

Setting **value** to **true** will enable automatic directory creation. Setting **value** to **false** will disable it.

## 10.4. CONFIGURING MESSAGE JOURNAL ATTRIBUTES

The attributes listed below are all child properties of the messaging server. Therefore, the command syntax for getting and setting their values using the management CLI is the same for each.

To read the current value of a given attribute, the syntax is as follows:

```
/subsystem=messaging-activemq/server=default:read-attribute(name=<ATTRIBUTE_NAME>)
```

The syntax for writing an attribute's value follows a corresponding pattern.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=
<ATTRIBUTE_NAME>,value=<NEW_VALUE>)
```

- **create-journal-dir**  
If this is set to true, the journal directory will be automatically created at the location specified in **journal-directory** if it does not already exist. The default value is **true**.
- **journal-buffer-timeout**  
Instead of flushing on every write that requires a flush, we maintain an internal buffer, and flush the entire buffer either when it is full, or when a timeout expires, whichever is sooner. This is used for both NIO and ASYNCIO and allows the system to scale better with many concurrent writes that require flushing.

This parameter controls the timeout at which the buffer will be flushed if it has not filled already. ASYNCIO can typically cope with a higher flush rate than NIO, so the system maintains different defaults for both NIO and ASYNCIO. The default for NIO is **3333333** nanoseconds, or 300 times per second. The default for ASYNCIO is **500000** nanoseconds, or 2000 times per second.



#### NOTE

By increasing the timeout, you may be able to increase system throughput at the expense of latency, the default parameters are chosen to give a reasonable balance between throughput and latency.

- **journal-buffer-size**  
The size, in bytes, of the timed buffer on ASYNCIO. Both **journal-buffer-size** and **journal-file-size** must be set larger than **min-large-message-size**. Otherwise, messages will not be written to the journal. See [Configuring Large Messages](#) for more information.
- **journal-compact-min-files**  
The minimal number of files before we can consider compacting the journal. The compacting algorithm won't start until you have at least **journal-compact-min-files**.

Setting this to **0** will disable the feature to compact completely. This could be dangerous though as the journal could grow indefinitely. Use it wisely!

The default for this parameter is **10**

- **journal-compact-percentage**  
The threshold to start compacting. When less than this percentage is considered live data, we start compacting. Note also that compacting will not kick in until you have at least **journal-compact-min-files** data files on the journal

The default for this parameter is **30**.

- **journal-file-size**  
The size of each journal file, in bytes. The default value for this is **10485760** bytes, or 10MB. Both **journal-file-size** and **journal-buffer-size** must be set larger than **min-large-message-size**. Otherwise, messages will not be written to the journal. See [Configuring Large Messages](#) for more information.
- **journal-max-io**

Write requests are queued up before being submitted to the system for execution. This parameter controls the maximum number of write requests that can be in the IO queue at any one time. If the queue becomes full then writes will block until space is freed up.

The system maintains different defaults for this parameter depending on whether it's NIO or ASYNCIO. The default for NIO is **1**, and the default for ASYNCIO is **500**.

There is a limit and the total max ASYNCIO cannot be higher than what is configured at the OS level, found at `/proc/sys/fs/aio-max-nr`, usually **65536**.

- **journal-min-files**

The minimum number of files the journal will maintain. When JBoss EAP starts and there is no initial message data, JBoss EAP will pre-create **journal-min-files** number of files. The default is **2**.

Creating journal files and filling them with padding is a fairly expensive operation and we want to minimize doing this at run-time as files get filled. By pre-creating files, as one is filled the journal can immediately resume with the next one without pausing to create it.

Depending on how much data you expect your queues to contain at steady state you should tune this number of files to match that total amount of data.

- **journal-pool-files**

The number of journal files that can be reused. ActiveMQ will create as many files as needed however when reclaiming files it will shrink back to the value. The default is **-1**, which means no limit.

- **journal-sync-transactional**

If this is set to true then JBoss EAP will make sure all transaction data is flushed to disk on transaction boundaries, such as a commit, prepare, or rollback. The default value is **true**.

- **journal-sync-non-transactional**

If this is set to true then JBoss EAP will make sure non transactional message data, such as sends and acknowledgements, are flushed to disk each time. The default value is **true**.

- **journal-type**

Valid values are **NIO** or **ASYNCIO**.

Choosing **NIO** tells JBoss EAP to use a Java NIO journal. **ASYNCIO** tells it to use a Linux asynchronous IO journal. If you choose **ASYNCIO** but are not running Linux, or you do not have `libaio` installed, JBoss EAP will use a Java NIO journal.

## 10.5. NOTE ON DISABLING DISK WRITE CACHE

This happens irrespective of whether you have executed a `fsync()` from the operating system or correctly synced data from inside a Java program!

By default many systems ship with disk write cache enabled. This means that even after syncing from the operating system there is no guarantee the data has actually made it to disk, so if a failure occurs, critical data can be lost.

Some more expensive disks have non volatile or battery backed write caches which won't necessarily lose data on event of failure, but you need to test them!



If your disk does not have an expensive non volatile or battery backed cache and it's not part of some kind of redundant array (e.g. RAID), and you value your data integrity you need to make sure disk write cache is disabled.

Be aware that disabling disk write cache can give you a nasty shock performance wise. If you've been used to using disks with write cache enabled in their default setting, unaware that your data integrity could be compromised, then disabling it will give you an idea of how fast your disk can perform when acting really reliably.

On Linux you can inspect or change your disk's write cache settings using the tools **hdparm** (for IDE disks) or **sdparm** or **sginfo** (for SDSI/SATA disks)

On Windows you can check / change the setting by right clicking on the disk and clicking properties.

## 10.6. INSTALLING LIBAIO

The Java NIO journal is highly performant, but if you are running JBoss EAP messaging using Linux Kernel 2.6 or later, Red Hat highly recommends that you use the ASYNCIO journal for the very best persistence performance.



### NOTE

JBoss EAP supports ASYNCIO only when installed on versions 6 or 7 of Red Hat Enterprise Linux and only when using the ext4, xfs, gfs2 or nfs4 file systems. It is not possible to use the ASYNCIO journal under other operating systems or earlier versions of the Linux kernel.

You will need **libaio** installed to use the ASYNCIO journal. You can easily install it using the following YUM command:

```
yum install libaio
```



### WARNING

Do not place your messaging journals on a tmpfs file system, which is used for the **/tmp** directory for example. JBoss EAP will fail to start if the ASYNCIO journal is using tmpfs.

## 10.7. CONFIGURING THE NFS SHARED STORE FOR MESSAGING

When using dedicated, shared store, high availability for data replication, you must configure both the live server and the backup server to use a shared directory on the NFS client. If you configure one server to use a shared directory on the NFS server and the other server to use a shared directory on the NFS client, the backup server cannot recognize when the live server starts or is running. So to work properly, both servers must specify a shared directory on the NFS client.

You must also configure the following options for the NFS client mount:

- **sync**: This option specifies that all changes are immediately flushed to disk.

- **intr**: This option allows NFS requests to be interrupted if the server goes down or cannot be reached.
- **noac**: This option disables attribute caching and is needed to achieve attribute cache coherence among multiple clients.
- **soft**: This option specifies that if the host serving the exported file system is unavailable, the error should be reported rather than waiting for the server to come back online.
- **lookupcache=none**: This option disables lookup caching.


See the [Shared Store](#) section in this guide for more information about how to use a shared file system for high availability.

## 10.8. CONFIGURING JBOSS EAP MESSAGING FOR ZERO PERSISTENCE

In some situations, zero persistence is required for a messaging system. Zero persistence means that no bindings data, message data, large message data, duplicate id caches, or paging data should be persisted.

To configure the **messaging-activemq** subsystem to perform zero persistence, set the **persistence-enabled** parameter to **false**.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=persistence-enabled,value=false)
```



### IMPORTANT

Be aware that if persistence is disabled, but paging is enabled, page files continue to be stored in the location specified by the **paging-directory** element. Paging is enabled when the **address-full-policy** attribute is set to **PAGE**. If full zero persistence is required, be sure to configure the **address-full-policy** attribute of the **address-setting** element to use **BLOCK**, **DROP** or **FAIL**.

## 10.9. IMPORTING AND EXPORTING JOURNAL DATA

See the JBoss EAP 7 [Migration Guide](#) for information on importing and exporting journal data.

## CHAPTER 11. CONFIGURING PAGING

### 11.1. ABOUT PAGING

JBoss EAP messaging supports many message queues with each queue containing millions of messages. The JBoss EAP messaging server runs with limited memory thereby making it difficult to store all message queues in memory at one time.

Paging is a mechanism used by the JBoss EAP messaging server to transparently page messages in and out of memory on an as-needed basis in order to accommodate large message queues in a limited memory.

JBoss EAP messaging starts paging messages to disk, when the size of messages in memory for a particular address exceeds the maximum configured message size.



#### NOTE

JBoss EAP messaging paging is enabled by default.

### 11.2. PAGE FILES

There is an individual folder for each address on the file system which stores messages in multiple files. These files which store the messages are called page files. Each file contains messages up to the maximum configured message size set by the **page-size-bytes** attribute.

The system navigates the page files as needed and removes the page files as soon as all messages in the page were received by client.



#### WARNING

For performance reasons, JBoss EAP messaging does not scan paged messages. Therefore, you should disable paging on a queue that is configured to group messages or to provide a last value. Also, message prioritization and message selectors will not behave as expected for queues that have paging enabled. You must disable paging for these features to work as expected.

For example, if a consumer has a message selector to read messages from a queue, only the messages in memory that match the selector are delivered to the consumer. When the consumer acknowledges delivery of these messages, new messages are de-paged and loaded into memory. There may be messages that match a consumer's selector on disk in page files but JBoss EAP messaging does not load them into memory until another consumer reads the messages in memory and provides free space. If the free space is not available, the consumer employing a selector may not receive any new messages.

### 11.3. CONFIGURING THE PAGING DIRECTORY

You can read the configuration for the paging directory by using the management CLI command below. In this example, the output displays the default configuration.

```

/subsystem=messaging-activemq/server=default/path=paging-directory:read-
resource
{
  "outcome" => "success",
  "result" => {
    "path" => "activemq/paging",
    "relative-to" => "jboss.server.data.dir"
  }
}

```

The **paging-directory** configuration element specifies the location on the file system to store the page files. JBoss EAP creates one folder for each paging address in this paging directory and the page files are stored within these folders. By default, this path is **activemq/paging/**. You can change the path location by using the following management CLI command.

```

/subsystem=messaging-activemq/server=default/path=paging-directory:write-
attribute(name=path,value=<PATH_LOCATION>)

```

Also note the **relative-to** attribute in the example output above. When **relative-to** is specified, the value of the **path** attribute is treated as relative to the file path specified by the **relative-to** attribute. By default, this value is the JBoss EAP **jboss.server.data.dir** property. For standalone servers, **jboss.server.data.dir** is located at **EAP\_HOME/standalone/data/**. For managed domains, each server will have its own **serverX/data/activemq/** directory located under **EAP\_HOME/domain/servers/**. You can change the value of **relative-to** using the following management CLI command.


```

/subsystem=messaging-activemq/server=default/path=paging-directory:write-
attribute(name=relative-to,value=<RELATIVE_LOCATION>)

```

## 11.4. CONFIGURING PAGING MODE

When messages delivered to an address exceed the configured size, that address goes into "paging mode".



### NOTE

Paging is done individually per address. If you configure a **max-size-bytes** for an address, it means each matching address will have a maximum size that you specified. However it does not mean that the total overall size of all matching addresses is limited to **max-size-bytes**.

Even with **page** mode, the server may crash due to an out-of-memory error. JBoss EAP messaging keeps a reference to each page file on the disk. In a situation with millions of page files, JBoss EAP messaging can face memory exhaustion. To minimize this risk, it is important to set the attribute **page-size-bytes** to a suitable value. You must configure the memory for your JBoss EAP messaging server to be greater than two times the number of destinations times the **max-size-bytes**, otherwise an out-of-memory error can occur.

You can read the current maximum size in bytes (**max-size-bytes**) for an address by using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/address-setting=
<ADDRESS_SETTING>:read-attribute(name=max-size-bytes)
```

You can configure the maximum size in bytes (**max-size-bytes**) for an address by using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/address-setting=
<ADDRESS_SETTING>:write-attribute(name=max-size-bytes,value=MAX_SIZE)
```

Use a similar syntax when reading or writing the values for the other paging-related attributes of an address setting. The table below lists each attribute, along with a description and a default value.

The following table describes the parameters on the address settings:

**Table 11.1. Paging Configuration for Address Settings**

Element	Description
address-full-policy	<p>This value of this attribute is used for paging decisions. The valid valid values are listed below.</p> <p><b>PAGE</b> Enables paging and page messages beyond the set limit to disk.</p> <p><b>DROP</b> Silently drops messages that exceed the set limit.</p> <p><b>FAIL</b> Drops messages and sends an exception to client message producers.</p> <p><b>BLOCK</b> Blocks client message producers when they send messages beyond the set limit.</p> <p>The default is <b>PAGE</b>.</p>
max-size-bytes	<p>This is used to specify the maximum memory size the address can have before entering into paging mode. The default is <b>10485760</b>.</p>
page-max-cache-size	<p>The system will keep page files up to <b>page-max-cache-size</b> in memory to optimize Input/Output during paging navigation. The default is <b>5</b>.</p>
page-size-bytes	<p>This is used to specify the size of each page file used on the paging system. The default is <b>2097152</b>.</p>



## IMPORTANT

By default, all addresses are configured to page messages after an address reaches **max-size-bytes**. If you do not want to page messages when the maximum size is reached, you can configure an address to drop messages, drop messages with an exception on client side, or block producers from sending further messages by setting the **address-full-policy** to **DROP**, **FAIL** and **BLOCK** respectively.

Be aware that if you change the **address-full-policy** from **PAGE** to **BLOCK** after any destination has started to page messages, consumers will no longer be able to consume paged messages.

## Addresses with Multiple Queues

When a message is routed to an address that has multiple queues bound to it, there is only a single copy of the message in memory. Each queue only handles a reference to this original copy of the message, so the memory is freed up only when all the queues referencing the original message, have delivered the message.



## NOTE

A single lazy queue/subscription can reduce the Input/Output performance of the entire address as all the queues will have messages being sent through an extra storage on the paging system.

## CHAPTER 12. WORKING WITH LARGE MESSAGES

JBoss EAP messaging supports large messages, even when the client or server has limited amounts of memory. Large messages can be streamed as they are, or they can be compressed further for more efficient transferral. A user can send a large message by setting an **InputStream** in the body of the message. When the message is sent, JBoss EAP messaging reads this **InputStream** and transmits data to the server in fragments.

Neither the client nor the server stores the complete body of a large message in memory. The consumer initially receives a large message with an empty body and thereafter sets an **OutputStream** on the message to stream it in fragments to a disk file.



### WARNING

When processing large messages, the server does not handle message properties in the same way as the message body. For example a message with a property set to a string that is bigger than **journal-buffer-size** cannot be processed by the server because it overfills the journal buffer.

### 12.1. STREAMING LARGE MESSAGES

JBoss EAP messaging supports setting the body of messages using the [java.io.InputStream](#) and [java.io.OutputStream](#) classes. Input streams are used directly for sending messages and output streams are used for receiving messages.

When receiving messages, there are two ways to deal with the output stream:

- You can block while the output stream is recovered using the **ClientMessage.saveToOutputStream(OutputStream out)** method.
- You can use the **ClientMessage.setOutputStream(OutputStream out)** method to asynchronously write the message to the stream. This method requires that the consumer be kept alive until the message has been fully received.

You can use any kind of stream you like, for example files, JDBC Blobs, or SocketInputStream, as long as it implements [java.io.InputStream](#) for sending messages and [java.io.OutputStream](#) for receiving messages.

#### Streaming Large Messages Using the Core API

The following table shows the methods available on the [ClientMessage](#) class that are available through JMS by using object properties.

ClientMessage Method	Description	JMS Equivalent Property
<b>setBodyInputStream(InputStream)</b>	Set the <b>InputStream</b> used to read a message body when it is sent.	<b>JMS_AMQ_InputStream</b>

ClientMessage Method	Description	JMS Equivalent Property
<b>setOutputStream(OutputStream)</b>	Set the <b>OutputStream</b> that will receive the body of a message. This method does not block.	<b>JMS_AMQ_OutputStream</b>
<b>saveOutputStream(OutputStream)</b>	Save the body of the message to the <b>OutputStream</b> . It will block until the entire content is transferred to the <b>OutputStream</b> .	<b>JMS_AMQ_SaveStream</b>

The following code example sets the output stream when receiving a core message.

```
ClientMessage firstMessage = consumer.receive(...);

// Block until the stream is transferred
firstMessage.saveOutputStream(firstOutputStream);

ClientMessage secondMessage = consumer.receive(...);

// Do not wait for the transfer to finish
secondMessage.setOutputStream(secondOutputStream);
```

The following code example sets the input stream when sending a core message:

```
ClientMessage clientMessage = session.createMessage();
clientMessage.setInputStream(dataInputStream);
```



#### NOTE

For messages larger than 2GiB, you must use the **\_AMQ\_LARGE\_SIZE** message property. This is because the **getBodySize()** method will return an invalid value because it is limited to the maximum integer value.

### Streaming Large Messages Over JMS

When using JMS, JBoss EAP messaging maps the core API streaming methods by setting object properties. You use the **Message.setObjectProperty(String name, Object value)** method to set the input and output streams.

The **InputStream** is set using the **JMS\_AMQ\_InputStream** property on messages being sent.

```
BytesMessage bytesMessage = session.createBytesMessage();
FileInputStream fileInputStream = new FileInputStream(fileInput);
BufferedInputStream bufferedInput = new
BufferedInputStream(fileInputStream);
bytesMessage.setObjectProperty("JMS_AMQ_InputStream", bufferedInput);
someProducer.send(bytesMessage);
```

The **OutputStream** is set using the **JMS\_AMQ\_SaveStream** property on messages being received in a blocking manner.



```

BytesMessage messageReceived = (BytesMessage)
messageConsumer.receive(120000);
File outputFile = new File("huge_message_received.dat");
FileOutputStream fileOutputStream = new FileOutputStream(outputFile);
BufferedOutputStream bufferedOutput = new
BufferedOutputStream(fileOutputStream);

// This will block until the entire content is saved on disk
messageReceived.setObjectProperty("JMS_AMQ_SaveStream", bufferedOutput);

```

The **OutputStream** can also be set in a non-blocking manner by using the **JMS\_AMQ\_OutputStream** property.

```

// This does not wait for the stream to finish. You must keep the consumer
active.
messageReceived.setObjectProperty("JMS_AMQ_OutputStream", bufferedOutput);

```



## NOTE

When streaming large messages using JMS, only **StreamMessage** and **BytesMessage** objects are supported.

## 12.2. CONFIGURING LARGE MESSAGES

### 12.2.1. Configure Large Message Location

You can read the configuration for the large messages directory by using the management CLI command below. The output is also included to highlight default configuration.

```

/subsystem=messaging-activemq/server=default/path=large-messages-
directory:read-resource
{
  "outcome" => "success",
  "result" => {
    "path" => "activemq/largemessages",
    "relative-to" => "jboss.server.data.dir"
  }
}

```



## IMPORTANT

To achieve the best performance, it is recommended to store the large messages directory on a different physical volume from the message journal or the paging directory.

The **large-messages-directory** configuration element is used to specify a location on the filesystem to store the large messages. Note that by default the path is **activemq/largemessages**. You can change the location for path by using the following management CLI command.

```

/subsystem=messaging-activemq/server=default/path=large-messages-
directory:write-attribute(name=path,value=<PATH_LOCATION>)

```

Also note the **relative-to** attribute in the output above. When **relative-to** is used, the value of the path attribute is treated as relative to the file path specified by **relative-to**. By default this value is the JBoss EAP property **jboss.server.data.dir**. For standalone servers, **jboss.server.data.dir** is located at **EAP\_HOME/standalone/data**. For domains, each server will have its own **serverX/data/activemq** directory located under **EAP\_HOME/domain/servers**. You can change the value of **relative-to** using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/path=large-messages-
directory:write-attribute(name=relative-to,value=<RELATIVE_LOCATION>)
```

### Configuring Large Message Size

Use the management CLI to view the current configuration for large messages. Note that this configuration is part of a **connection-factory** element. For example, to read the current configuration for the default RemoteConnectionFactory that is included, use the following command:

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:read-attribute(name=min-large-message-
size)
```

Set the attribute using a similar syntax.

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=min-large-message-
size,value=<NEW_MIN_SIZE>)
```



#### NOTE

The value of the attribute **min-large-message-size** should be in bytes.

### Configuring Large Message Compression

You can choose to compress large messages for fast and efficient transfer. All compression/decompression operations are handled on the client side. If the compressed message is smaller than **min-large-message size**, it is sent to the server as a regular message. Compress large messages by setting the boolean property **compress-large-messages** to **true** using the management CLI.

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=compress-large-
messages,value=true)
```

## 12.2.2. Configuring Large Message Size Using the Core API

If you are using the core API on the client side, you need to use the **setMinLargeMessageSize** method to specify the minimum size of large messages. The minimum size of large messages (**min-large-message-size**) is set to 100KB by default.

```
ServerLocator locator = ActiveMQClient.createServerLocatorWithoutHA(new
TransportConfiguration(InVMConnectorFactory.class.getName()))

locator.setMinLargeMessageSize(25 * 1024);
```

```
ClientSessionFactory factory =  
ActiveMQClient.createClientSessionFactory();
```

## CHAPTER 13. SCHEDULING MESSAGES

You can specify a time in the future, at the earliest, for a message to be delivered. This can be done by setting the `_AMQ_SCHED_DELIVERY` scheduled delivery property before the message is sent.

The specified value must be a positive **long** that corresponds to the time in milliseconds for the message to be delivered. Below is an example of sending a scheduled message using the JMS API.

```
// Create a message to be delivered in 5 seconds
TextMessage message = session.createTextMessage("This is a scheduled
message message that will be delivered in 5 sec.");
message.setLongProperty("_AMQ_SCHED_DELIVERY", System.currentTimeMillis()
+ 5000);
producer.send(message);

...

// The message will not be received immediately, but 5 seconds later
TextMessage messageReceived = (TextMessage) consumer.receive();
```

Scheduled messages can also be sent using the core API by setting the `_AMQ_SCHED_DELIVERY` property before sending the message.

## CHAPTER 14. TEMPORARY QUEUES AND RUNTIME QUEUES

When designing a request-reply pattern where a client sends a request and waits for a reply, you must consider whether each runtime instance of the client requires a dedicated queue for its replies, or whether the runtime instances can access a shared queue, selecting their specific reply messages based on an appropriate attribute.

If multiple queues are required, then clients need the ability to create a queue dynamically. JMS provides this facility using the concept of temporary queues. A **TemporaryQueue** is created on request by the **Session**. It exists for the life of the **Connection**, for example until the connection is closed, or until the temporary queue is deleted. This means that although the temporary queue is created by a specific session, it can be reused by any other sessions created from the same connection.

The trade-off between using a shared queue and individual temporary queues for replies is influenced by the potential number of active client instances. With a shared-queue approach, at some provider-specific threshold, contention for access to the queue can become a concern. This has to be contrasted against the additional overhead associated with the provider creating queue storage at runtime and the impact on machine memory of hosting a potentially large number of temporary queues.

The following example creates a temporary queue and consumer for each client on startup. It sets the **JMSReplyTo** property on each message to the temporary queue, and then sets a correlation ID on each message to correlate request messages to response messages. This avoids the overhead of creating and closing a consumer for each request, which is expensive. The same producer and consumer can be shared or pooled across many threads. Any messages that have been received, but not yet acknowledged when the session terminates, are retained and redelivered when a consumer next accesses the queue.

### Temporary Queue Code Example

```

...
// Create a temporary queue, one per client
Destination temporaryQueue = session.createTemporaryQueue();
MessageConsumer responseConsumer = session.createConsumer(temporaryQueue);

// This class handles messages to the temporary queue
responseConsumer.setMessageListener(this);

// Create the message to send
TextMessage textMessage = session.createTextMessage();
textMessage.setText("My new message!");

// Set the reply to field and correlation ID
textMessage.setJMSReplyTo(temporaryQueue);
textMessage.setJMSCorrelationID(myCorrelationID);

producer.send(textMessage);
...

```

In a similar manner, temporary topics are created using the **Session.createTemporaryTopic()** method.

## CHAPTER 15. FILTER EXPRESSIONS AND MESSAGE SELECTORS

The **messaging-activemq** subsystem in JBoss EAP provides a powerful filter language based on a subset of the SQL 92 expression syntax.

It is the same as the syntax used for JMS selectors, but the predefined identifiers are different. For documentation on JMS selector syntax, refer to the JEE 7 Javadoc for [javax.jms.Message](#).

The **filter** attribute can be found in several places within the configuration.

- **Predefined Queues.** When pre-defining a queue, a filter expression can be defined for it. Only messages that match the filter expression will enter the queue. The configuration snippet below shows a queue definition that includes a filter:

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  ...
  <queue
    name="myQueue"
    filter="<FILTER_EXPRESSION>"
    ...
  />
  ...
</subsystem>
```

To create queue with a selector in the management CLI you would use something like following command:

```
jms-queue add --queue-address=<QUEUE_ADDRESS> --selector=
<FILTER_EXPRESSION>
```

- **Core bridges** can be defined with an optional filter expression, only matching messages will be bridged. Below is a snippet from a sample configuration file where the **messaging-activemq** subsystem includes a bridge with a filter.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  ...
  <bridge
    name="myBridge"
    filter="<FILTER_EXPRESSION>"
    ...
  />
  ...
</subsystem>
```

- **Diverts** can be defined with an optional filter expression, only matching messages will be diverted. See [Diverts](#) for more information. The example snippet below shows a Divert using a filter:

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  ...
  <divert
    name="myDivert"
    filter="<FILTER_EXPRESSION>"
  />
  ...
</subsystem>
```

```
    ...  
  />  
  ...  
</subsystem>
```

There are some differences between JMS selector expressions and JBoss EAP messaging core filter expressions. Whereas JMS selector expressions operate on a JMS message, JBoss EAP messaging core filter expressions operate on a core message.

The following identifiers can be used in core filter expressions to refer to the attributes of a core message:

- **AMQPriority**. To refer to the priority of a message. Message priorities are integers with valid values from 0 - 9. 0 is the lowest priority and 9 is the highest. E.g. `AMQPriority = 3 AND animal = 'aardvark'`
- **AMQExpiration**. To refer to the expiration time of a message. The value is a long integer.
- **AMQDurable**. To refer to whether a message is durable or not. The value is a string with valid values: `DURABLE` or `NON_DURABLE`.
- **AMQTimestamp**. The timestamp of when the message was created. The value is a long integer.
- **AMQSize**. The size of a message in bytes. The value is an integer.

Any other identifiers used in core filter expressions will be assumed to be properties of the message.

## CHAPTER 16. CONFIGURING MESSAGE EXPIRY

Sent messages can be set to expire on the server if they are not delivered to a consumer after a specified amount of time. These expired messages can later be consumed for further inspection.

### Set Message Expiry Using the Core API

Using the core API, you can set an expiration time on a message using the `setExpiration` method.

```
// The message will expire 5 seconds from now
message.setExpiration(System.currentTimeMillis() + 5000);
```

### Set Message Expiry Using JMS

You can set the time to live for the JMS `MessageProducer` to use when sending messages. You specify this value, in milliseconds, using the `setTimeToLive` method.

```
// Messages sent by this producer will be retained for 5 seconds before
expiring
producer.setTimeToLive(5000);
```

You can also specify the message expiry on a per-message basis by setting the time to live on the producer's `send` method.

```
// The last parameter of the send method is the time to live, in
milliseconds
producer.send(message, DeliveryMode.PERSISTENT, 0, 5000)
```

Expired messages that are consumed from an expiry address have the following properties.

- `_AMQ_ORIG_ADDRESS`  
A `String` property containing the original address of the expired message.
- `_AMQ_ACTUAL_EXPIRY`  
A `Long` property containing the actual expiration time of the expired message.

### 16.1. EXPIRY ADDRESS

You can specify where to send expired messages by setting an expiry address. If a message expires and no expiry address is specified, the message is removed from the queue and dropped.

You can set an `expiry-address` for an `address-setting` using the management CLI. In the below example, expired messages in the `jms.queue.exampleQueue` queue will be sent to the `jms.queue.expiryQueue` expiry address.

```
/subsystem=messaging-activemq/server=default/address-
setting=jms.queue.exampleQueue:write-attribute(name=expiry-
address,value=jms.queue.expiryQueue)
```

### 16.2. EXPIRY REAPER THREAD

A reaper thread periodically inspects the queues to check whether messages have expired. You can set the scan period and thread priority for the reaper thread using the management CLI.



Set the scan period for the expiry reaper thread, which is how often, in milliseconds, the queues will be scanned to detect expired messages. The default is **30000**. You can set this to **-1** to disable the reaper thread.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=message-expiry-scan-period,value=30000)
```

Set the thread priority for the expiry reaper thread. Possible values are from **0** to **9**, with **9** being the highest priority. The default is **3**.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=message-expiry-thread-priority,value=3)
```

## CHAPTER 17. CONFIGURING DELAYED REDELIVERY

Delayed redelivery to an address is defined by the **redelivery-delay** attribute of an **address-setting** configuration element. If a redelivery delay is specified, JBoss EAP waits for the duration of this delay before redelivering messages. If **redelivery-delay** is set to **0**, there is no redelivery delay. To get the current value of **redelivery-delay** for a given **address-setting**, use the following management CLI command as an example.

```
/subsystem=messaging-activemq/server=default/address-setting=
<YOUR_ADDRESS_SETTING>:read-attribute(name=redelivery-delay)
```

The table below lists the configuration attributes of an **address-setting** that can be used to configure the redelivery of messages. Set the value for a given attribute using the following management CLI command as an example.

```
/subsystem=messaging-activemq/server=default/address-setting=
<YOUR_ADDRESS_SETTING>:write-attribute(name=<ATTRIBUTE>, value=<NEW_VALUE>)
```

**Table 17.1. Delivery Related Attributes of Address Settings**

Attribute	Description
max-delivery-attempts	Defines how many time a canceled message can be redelivered before sending to the <b>dead-letter-address</b> . The default is <b>10</b> .
max-redelivery-delay	Maximum value for the <b>redelivery-delay</b> (in ms). The default is <b>0</b> .
redelivery-delay	Defines how long to wait in milliseconds before attempting redelivery of a canceled message. The default is <b>0</b> .
redelivery-multiplier	Multiplier to apply to the <b>redelivery-delay</b> parameter. The default is <b>1.0</b>

See [Address Settings](#) for details on configuring an **address-setting**.

## CHAPTER 18. CONFIGURING DEAD LETTER ADDRESSES

A dead letter address is defined in the **address-setting** element of the **messaging-activemq** subsystem configuration. To read the current configuration for a given **address-setting**, use the following management CLI command as an example.

```
/subsystem=messaging-activemq/server=default/address-setting=  
<ADDRESS_SETTING>:read-attribute(name=dead-letter-address)
```

If a **dead-letter-address** is not specified, messages are removed after trying to deliver **max-delivery-attempts** times. By default, messages delivery is attempted 10 times. Setting **max-delivery-attempts** to **-1** allows infinite redelivery attempts. The example management CLI commands below illustrate how to set the **dead-letter-address** and the **max-delivery-attempts** attributes for a given **address-setting**.

```
/subsystem=messaging-activemq/server=default/address-setting=  
<ADDRESS_SETTING>:write-attribute(name=dead-letter-address,value=  
<NEW_VALUE>)
```

```
/subsystem=messaging-activemq/server=default/address-setting=  
<ADDRESS_SETTING>:write-attribute(name=max-delivery-attempts,value=  
<NEW_VALUE>)
```

For example, a dead letter can be set globally for a set of matching addresses and you can set **max-delivery-attempts** to **-1** for a specific address setting to allow infinite redelivery attempts only for this address. Address wildcards can also be used to configure dead letter settings for a set of addresses.

See [Address Settings](#) for details on creating and configuring an **address-setting**.

## CHAPTER 19. FLOW CONTROL

Flow control can be used to limit the flow of messaging data between a client and server so that messaging participants are not overwhelmed. You can manage the flow of data from both the consumer side and the producer side.

### 19.1. CONSUMER FLOW CONTROL

JBoss EAP messaging includes configuration that defines how much data to pre-fetch on behalf of consumers and that controls the rate at which consumers can consume messages.

#### Window-based flow control

JBoss EAP messaging pre-fetches messages into a buffer on each consumer. The size of the buffer is determined by the **consumer-window-size** attribute of a **connection-factory**. The example configuration below shows a **connection-factory** with the **consumer-window-size** attribute explicitly set.

```
<connection-factory name="MyConnFactory" ... consumer-window-size="1048576" />
```

Use the management CLI to read and write the value of **consumer-window-size** attribute for a given **connection-factory**. The examples below show how this done using the **InVmConnectionFactory** connection factory, which is the default for consumers residing in the same virtual machine as the server, e.g. a local **MessageDrivenBean**.

- Read the **consumer-window-size** attribute of the **InVmConnectionFactory** from the management CLI

```
/subsystem=messaging-activemq/server=default/connection-factory=InVmConnectionFactory:read-attribute(name=consumer-window-size)
{
  "outcome" => "success",
  "result" => 1048576
}
```

- Write the **consumer-window-size** attribute from the management CLI

```
/subsystem=messaging-activemq/server=default/connection-factory=InVmConnectionFactory:write-attribute(name=consumer-window-size,value=1048576)
{"outcome" => "success"}
```

The value for **consumer-window-size** must be an integer. Some values have special meaning as noted in the table below.

**Table 19.1. Values for consumer-window-size**

Value	Description
-------	-------------

Value	Description
n	An integer value used to set the buffer's size to <b>n</b> bytes. The default is <b>1048576</b> , which should be fine in most cases. Benchmarking will help you find an optimal value for the window size if the default value is not adequate.
0	Turns off buffering. This can help with slow consumers and can give deterministic distribution across multiple consumers.
-1	Creates an unbounded buffer. This can help facilitate very fast consumers that pull and process messages as quickly as they are received.



### WARNING

Setting **consumer-window-size** to **-1** can overflow the client memory if the consumer is not able to process messages as fast as it receives them.

If you are using the core API, the consumer window size can be set from the **ServerLocator** using its **setConsumerWindowSize()** method.

If you are using JMS, the client can specify the consumer window size by using the **setConsumerWindowSize()** method of the instantiated **ConnectionFactory**.

### Rate-limited flow control

JBoss EAP messaging can regulate the rate of messages consumed per second, a flow control method known as throttling. Use the **consumer-max-rate** attribute of the appropriate **connection-factory** to ensure that a consumer never consumes messages at a rate faster than specified.

```
<connection-factory name="MyConnFactory" ... consumer-max-rate="10" />
```

The default value is **-1**, which disables rate limited flow control.

The management CLI is the recommended way to read and write the **consumer-max-rate** attribute. The examples below show how this done using the **InVmConnectionFactory** connection factory, which is the default for consumers residing in the same virtual machine as the server, e.g. a local **MessageDrivenBean**.

- Read the **consumer-max-rate** attribute using the management CLI

```
/subsystem=messaging-activemq/server=default/connection-
factory=InVmConnectionFactory:read-attribute(name=consumer-max-rate)
{
    "outcome" => "success",
    "result" => -1
}
```

- Write the **consumer-max-rate** attribute using the management CLI:

```
/subsystem=messaging-activemq/server=default/connection-
factory=InVmConnectionFactory:write-attribute(name=consumer-max-
rate,value=100)
{"outcome" => "success"}
```

If you are using JMS the max rate size can be set using `setConsumerMaxRate(int consumerMaxRate)` method of the instantiated **ConnectionFactory**.

If you are using the Core API the rate can be set with the **ServerLocator.setConsumerMaxRate(int consumerMaxRate)** method.

## 19.2. PRODUCER FLOW CONTROL

JBoss EAP messaging can also limit the amount of data sent from a client in order to prevent the server from receiving too many messages.

### Window-based flow control

JBoss EAP messaging regulates message producers by using an exchange of credits. Producers can send messages to an address as long as they have sufficient credits to do so. The amount of credits required to send a message is determined by its size. As producers run low on credits, they must request more from the server. Within the server configuration, the amount of credits a producer can request at one time is known as the **producer-window-size**, an attribute of the **connection-factory** element:

```
<connection-factory name="MyConnFactory" ... producer-window-
size="1048576" />
```

The window size determines the amount of bytes that can be in-flight at any one time, thus preventing the remote connection from overloading the server.

Use the management CLI to read and write the **producer-window-size** attribute of a given connection factory. The examples below use the **RemoteConnectionFactory**, which is included in the default configuration and intended for use by remote clients.

- Read the **producer-window-size** attribute using the management CLI:

```
subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:read-attribute(name=producer-window-size)
{
  "outcome" => "success",
  "result" => 65536
}
```

- Write the **producer-window-size** attribute using the management CLI:

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=producer-window-
size,value=65536)
{"outcome" => "success"}
```


If you are using JMS, the client can call the `setProducerWindowSize(int producerWindowSize)` method of the `ConnectionFactory` to set the window size directly.

If you are using the core API, the window size can be set using the `setProducerWindowSize(int producerWindowSize)` method of the `ServerLocator`.

### Blocking producer window-based flow control

Typically, the messaging server always provides the same number of credits that was requested. However, it is possible to limit the number of credits sent by the server, which can prevent it from running out of memory due to producers sending more messages than can be handled at one time.

For example, if you have a JMS queue called `myqueue` and you set the maximum memory size to 10MB, the server will regulate the number of messages in the queue so that its size never exceeds 10MB. When the address gets full, producers will block on the client side until sufficient space is freed up on the address.



#### NOTE

Blocking producer flow control is an alternative approach to paging, which does not block producers but instead pages messages to storage. See [About Paging](#) for more information.

The `address-setting` configuration element contains the configuration for managing blocking producer flow control. An `address-setting` is used to apply a set of configuration to all queues registered to that address. See [Configuring Address Settings](#) for more information on how this is done.

For each `address-setting` requiring blocking producer flow control, you must include a value for the `max-size-bytes` attribute. The total memory for all queues bound to that address cannot exceed `max-size-bytes`. In the case of JMS topics, this means the total memory of all subscriptions in the topic cannot exceed `max-size-bytes`.

You must also set the `address-full-policy` attribute to `BLOCK` so the server knows that producers should be blocked if `max-size-bytes` is reached. Below is an example `address-setting` with both attributes set:

```
<address-setting ...
  name="myqueue"
  address-full-policy="BLOCK"
  max-size-bytes="100000" />
```

The above example would set the maximum size of the JMS queue "myqueue" to **100000** bytes. Producers will be blocked from sending to that address once it has reached its maximum size.

Use the management CLI to set these attributes, as in the examples below:

- Set `max-size-bytes` for a specified `address-setting`

```
/subsystem=messaging-activemq/server=default/address-
setting=myqueue:write-attribute(name=max-size-bytes,value=100000)
{"outcome" => "success"}
```

- Set `address-full-policy` for a specified `address-setting`

```
/subsystem=messaging-activemq/server=default/address-
```

```
setting=myqueue:write-attribute(name=address-full-policy,value=BLOCK)
{"outcome" => "success"}
```

### Rate-limited flow control

JBoss EAP messaging limits the number of messages a producer can send per second if you specify a **producer-max-rate** for the **connection-factory** it uses, as in the example below:

```
<connection-factory name="MyConnFactory" producer-max-rate="1000" />
```

The default value is **-1**, which disables rate limited flow control.

Use the management CLI to read and write the value for **producer-max-rate**. The examples below use the **RemoteConnectionFactory**, which is included in the default configuration and intended for use by remote clients.

- Read the value of the **producer-max-rate** attribute:

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:read-attribute(name=producer-max-rate)
{
  "outcome" => "success",
  "result" => -1
}
```

- Write the value of a **producer-max-rate** attribute:

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=producer-max-
rate,value=100)
{"outcome" => "success"}
```

If you use the core API, set the rate by using the method **ServerLocator.setProducerMaxRate(int producerMaxRate)**.

If you are using JNDI to instantiate and look up the connection factory, the max rate can be set on the client using the **setProducerMaxRate(int producerMaxRate)** method of the instantiated connection factory.



## CHAPTER 20. CONFIGURING PRE-ACKNOWLEDGMENTS

JMS specifies three acknowledgement modes:

- AUTO\_ACKNOWLEDGE
- CLIENT\_ACKNOWLEDGE
- DUPS\_OK\_ACKNOWLEDGE

In some cases you can afford to lose messages in the event of a failure, so it would make sense to acknowledge the message on the server before delivering it to the client. This extra mode is supported by JBoss EAP messaging and is called pre-acknowledge mode.

The disadvantage of pre-acknowledging on the server before delivery is that the message will be lost if the server's system crashes after acknowledging the message but before it is delivered to the client. In that case, the message is lost and will not be recovered when the system restarts.

Depending on your messaging case, pre-acknowledge mode can avoid extra network traffic and CPU usage at the cost of coping with message loss.

An example use case for pre-acknowledgement is for stock price update messages. With these messages, it might be reasonable to lose a message in event of a crash since the next price update message will arrive soon, overriding the previous price.



### NOTE

If you use pre-acknowledge mode, you will lose transactional semantics for messages being consumed since they are being acknowledged first on the server, not when you commit the transaction.

## 20.1. CONFIGURING THE SERVER

A connection factory can be configured to use pre-acknowledge mode by setting its **pre-acknowledge** attribute to **true** using the management CLI as below:

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=pre-
acknowledge,value=true)
```

## 20.2. CONFIGURING THE CLIENT

Pre-acknowledge mode can be configured in a client's JNDI context environment, e.g. in a **jndi.properties** file:

```
java.naming.factory.initial=org.apache.activemq.artemis.jndi.ActiveMQInitialContextFactory
connection.ConnectionFactory=tcp://localhost:8080?preAcknowledge=true
```

Alternatively, to use pre-acknowledge mode using the JMS API, create a JMS Session with the **ActiveMQSession.PRE\_ACKNOWLEDGE** constant.

```
// messages will be acknowledge on the server *before* being delivered to
```

*the client*

```
Session session = connection.createSession(false,  
ActiveMQJMSConstants.PRE_ACKNOWLEDGE);
```

## CHAPTER 21. INTERCEPTORS

JBoss EAP messaging supports interceptors to intercept packets entering and exiting the server. Incoming and outgoing interceptors are called for every packet entering or exiting the server respectively. This allows custom code to be executed, such as for auditing or filtering packets. Interceptors can modify the packets they intercept. This makes interceptors powerful, but also potentially dangerous.

### 21.1. IMPLEMENTING INTERCEPTORS

An interceptor must implement the `Interceptor` interface:

```
package org.apache.artemis.activemq.api.core.interceptor;

public interface Interceptor
{
    boolean intercept(Packet packet, RemotingConnection connection) throws
    ActiveMQException;
}
```

The returned boolean value is important:

- if true is returned, the process continues normally
- if false is returned, the process is aborted, no other interceptors will be called and the packet will not be processed further by the server.

Interceptor classes should be added to JBoss EAP as a module. See [Create a Custom Module](#) in the JBoss EAP [*Configuration Guide*] for more information.

### 21.2. CONFIGURING INTERCEPTORS

After adding their module to JBoss EAP as a custom module, both incoming and outgoing interceptors are added to the messaging subsystem configuration by using the management CLI.



#### NOTE

You must start JBoss EAP in "administrator only" mode before the new interceptor configuration will be accepted. See [Running JBoss EAP in Admin-Only Mode](#) in the JBoss EAP Configuration Guide] for details. Restart the server in normal mode after the new configuration is processed.

Each interceptor should be added according to the example management CLI command below. The examples assume each interceptor has already been added to JBoss EAP as a custom module.

```
/subsystem=messaging-activemq/server=default:list-add(name=incoming-
interceptors, value={name => "foo.bar.MyIncomingInterceptor",
module=>"foo.bar.interceptors"})
```

Adding an outgoing interceptor follows a similar syntax, as the example below illustrates.

```
/subsystem=messaging-activemq/server=default:list-add(name=outgoing-  
interceptors, value={name => "foo.bar.MyOutgoingInterceptor",  
module=>"foo.bar.interceptors"})
```

## CHAPTER 22. MESSAGE GROUPING

A message group is a group of messages that share certain characteristics:

- All messages in a message group are grouped under a common group id. This means that they can be identified with a common group property.
- All messages in a message group are serially processed and consumed by the same consumer, irrespective of the number of consumers on the queue. This means that a specific message group with a unique group id is always processed by one consumer when the consumer opens it. If the consumer closes the message group, then the entire message group is directed to another consumer in the queue.

Message groups are especially useful when there is a need for messages with a certain value of the property (group id) to be processed serially by a single consumer.



### IMPORTANT

Message grouping will not work as expected if the queue has paging enabled. Be sure to [disable paging](#) before configuring a queue for message grouping.

For information about configuring message grouping within a cluster of messaging servers, see [Clustered Message Grouping](#) in Part III, [Configuring Multiple Messaging Systems](#).

### 22.1. CONFIGURING MESSAGE GROUPS USING THE CORE API

The property `_AMQ_GROUP_ID` is used to identify a message group using the Core API on the client side. To pick a random unique message group identifier, you can also set the `auto-group` property to `true` on the `SessionFactory`.

### 22.2. CONFIGURING MESSAGE GROUPS USING JMS

The property `JMSXGroupID` is used to identify a message group for Java Messaging Service (JMS) clients. If you wish to send a message group with different messages to one consumer, you can set the same `JMSXGroupID` for different messages.

```
Message message = ...
message.setStringProperty("JMSXGroupID", "Group-0");
producer.send(message);

message = ...
message.setStringProperty("JMSXGroupID", "Group-0");
producer.send(message);
```

An alternative approach is to use the one of the following attributes of the `connection-factory` to be used by the client: `auto-group` or `group-id`.

When `auto-group` is set to `true`, the `connection-factory` will begin to use a random unique message group identifier for all messages sent through it. You can use the management CLI to set the `auto-group` attribute.

```
/subsystem=messaging-activemq/server=default/connection-  
factory=RemoteConnectionFactory:write-attribute(name=auto-  
group,value=true)
```

The **group-id** attribute will set the property **JMSXGroupID** to the specified value for all messages sent through the connection factory. To set a specific **group-id** on the connection factory, use the management CLI.

```
/subsystem=messaging-activemq/server=default/connection-  
factory=RemoteConnectionFactory:write-attribute(name=group-  
id,value="Group-0")
```

## CHAPTER 23. DIVERTS

Diverts are objects configured in JBoss EAP messaging that help in diverting messages from one address to another. Diverts can be classified into the following types:

### Exclusive

A message is only diverted to a new address and never sent to the old address.

### Non-exclusive

A message is sent the old address, and a copy of it is also sent to the new address. Non-exclusive diverts can be used for splitting the flow of messages.

A divert will only divert a message to an address on the same server. If you want to divert to an address on a different server, a common pattern would be to divert to a local store-and-forward queue, then set up a bridge that consumes from that queue and forwards to an address on a different server.

Diverts are therefore a very sophisticated concept. When combined with bridges, diverts can be used to create interesting and complex routings. The set of diverts on a server can be thought of as a type of routing table for messages. Combining diverts with bridges allows you to create a distributed network of reliable routing connections between multiple geographically distributed servers, creating your global messaging mesh. See [Configuring Core Bridges](#) for more information on how to use bridges.

Diverts can be configured to apply a **Transformer** and an optional message filter. An optional message filter helps to divert only messages which match the specified filter. For more information on filters see [Filter Expressions and Message Selectors](#).

A transformer is used for transforming messages to another form. When a transformer is specified, all diverted messages are transformed by the **Transformer**. All transformers must implement the `org.apache.activemq.artemis.core.server.cluster.Transformer` interface:

```
package org.apache.activemq.artemis.core.server.cluster;
import org.apache.activemq.artemis.core.server.ServerMessage;

public interface Transformer {
    ServerMessage transform(ServerMessage message);
}
```

To enable JBoss EAP messaging to instantiate an instance of your transformer implementation, it should be included in a JBoss EAP module, and the module should be added as a dependency to the `org.apache.activemq.artemis` module. See [Create a Custom Module](#) in the JBoss EAP *Configuration Guide* for information on how to create a custom module. To add a dependency to the `org.apache.activemq.artemis` module, open the file `EAP_HOME/modules/system/layers/base/org/apache/activemq/artemis/main/module.xml` in a text editor and add your `<module>` to the list of `<dependencies>` like so:

```
<module xmlns="urn:jboss:module:1.3" name="org.apache.activemq.artemis">
  <resources>
    ...
  </resources>
  <dependencies>
    ...
    <module name="<YOUR_MODULE_NAME>" />
  </dependencies>
</module>
```

## 23.1. EXCLUSIVE DIVERTS

Below is an example of an exclusive divert as it might appear in configuration:

```
<divert
  name="prices-divert"
  address="jms.topic.priceUpdates"
  forwarding-address="jms.queue.priceForwarding"
  filter="office='New York'"
  transformer-class-
name="org.apache.activemq.artemis.jms.example.AddForwardingTimeTransformer"
  exclusive="true"/>
```

In this example, a divert called **prices-divert** is configured that will divert any messages sent to the address **jms.topic.priceUpdates**, which maps to any messages sent to a JMS topic called **priceUpdates**, to another local address **jms.queue.priceForwarding**, corresponding to a local JMS queue called **priceForwarding**

We also specify a message filter string so that only messages with the message property **office** with a value of **New York** will be diverted. All other messages will continue to be routed to the normal address. The filter string is optional, if not specified then all messages will be considered matched.

Note that a transformer class is specified. In this example the transformer simply adds a header that records the time the divert happened.

This example is actually diverting messages to a local store and forward queue, which is configured with a bridge that forwards the message to an address on another server. See [Configuring Core Bridges](#) for more details.

## 23.2. NON-EXCLUSIVE DIVERTS

Below is an example of a non-exclusive divert. Non exclusive diverts can be configured in the same way as exclusive diverts with an optional filter and transformer.

```
<divert
  name="order-divert"
  address="jms.queue.orders"
  forwarding-address="jms.topic.spytopic"
  exclusive="false"/>
```

The above divert takes a copy of every message sent to the address **jms.queue.orders**, which maps to a JMS Queue called **orders**, and sends it to a local address called **jms.topic.SpyTopic**, corresponding to a JMS topic called **SpyTopic**).

### Creating diverts

Use the management CLI to create the type of divert you want:

```
/subsystem=messaging-activemq/server=default/divert=my-divert:add(divert-
address=news.in,forwarding-address=news.forward)
```

Non-exclusive diverts are created by default. To create an exclusive divert use the **exclusive** attribute:



```
/subsystem=messaging-activemq/server=default/divert=my-exclusive-
divert:add(divert-address=news.in,forwarding-
address=news.forward,exclusive=true)
```

The below table captures a divert's attributes and their description. You can have the management CLI display this information using the following command:

```
/subsystem=messaging-activemq/server=default/divert=*:read-resource-
description()
```

Attribute	Description
divert-address	Address to divert from. Required.
exclusive	Whether the divert is exclusive, meaning that the message is diverted to the new address, and does not go to the old address at all. The default is false.
filter	An optional filter string. If specified then only messages which match the filter expression will be diverted.
forwarding-address	Address to divert to. Required.
routing-name	Routing name of the divert.
transformer-class-name	The name of a class used to transform the message's body or properties before it is diverted.

## CHAPTER 24. THREAD MANAGEMENT

Each JBoss EAP messaging server maintains a single thread pool for general use, and scheduled thread pool for scheduled use. A Java scheduled thread pool cannot be configured to use a standard thread pool, otherwise we could use a single thread pool for both scheduled and non scheduled activity.

Note that JBoss EAP uses the new, non-blocking NIO. By default, JBoss EAP messaging uses a number of threads equal to three times the number of cores, or hyper-threads, as reported by `.getRuntime().availableProcessors()` for processing incoming packets. To override this value, set the number of threads by specifying the `nio-remoting-threads` parameter in the transport configuration. See [Configuring the Messaging Transports](#) for more information.

### 24.1. SERVER SCHEDULED THREAD POOL

The server scheduled thread pool is used for most activities on the server side that require running periodically or with delays. It maps internally to a `java.util.concurrent.ScheduledThreadPoolExecutor` instance.

The maximum number of thread used by this pool is configured using the `scheduled-thread-pool-max-size` parameter. The default value is 5 threads. A small number of threads is usually sufficient for this pool. To change this value for the default JBoss EAP messaging server, use the following management CLI command:

```
/subsystem=messaging-activemq/server=default:write-attribute(name=scheduled-thread-pool-max-size,value=10)
```

### 24.2. SERVER GENERAL PURPOSE THREAD POOL

The general purpose thread pool is used for most asynchronous actions on the server side. It maps internally to a `java.util.concurrent.ThreadPoolExecutor` instance.

The maximum number of threads used by this pool is configured using the `thread-pool-max-size` attribute.

If `thread-pool-max-size` is set to `-1`, the thread pool has no upper bound and new threads are created on demand if there are not enough threads available to fulfill a request. If activity later subsides, then threads are timed out and closed.

If `thread-pool-max-size` is set to a positive integer greater than zero, the thread pool is bounded. If requests come in and there are no free threads available in the pool, requests will block until a thread becomes available. It is recommended that a bounded thread pool be used with caution since it can lead to deadlock situations if the upper bound is configured too low.

The default value for `thread-pool-max-size` is `30`. To set a new value for the default JBoss EAP messaging server, use the following management CLI command.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=thread-pool-max-size,value=40)
```

See the [Java 8 Javadoc](#) for more information on unbounded (cached), and bounded (fixed) thread pools.

### 24.3. EXPIRY REAPER THREAD

A single thread is also used on the server side to scan for expired messages in queues. We cannot use either of the thread pools for this since this thread needs to run at its own configurable priority.

## 24.4. ASYNCHRONOUS IO

Asynchronous IO has a thread pool for receiving and dispatching events out of the native layer. It is on a thread dump with the prefix **ArtemisMQ-AIO-poller-pool**. JBoss EAP messaging uses one thread per opened file on the journal (there is usually one).

There is also a single thread used to invoke writes on libaio. It is done to avoid context switching on libaio that would cause performance issues. This thread is found on a thread dump with the prefix **ArtemisMQ-AIO-writer-pool**.

## 24.5. CLIENT THREAD MANAGEMENT

JBoss EAP includes a client thread pool used for creating client connections. This pool is separate from the static pools mentioned earlier in this chapter and is used by JBoss EAP when it behaves like a client. For example, client thread pool clients are created as cluster connections with other nodes in the same cluster, or when the Artemis resource adapter connects to a remote Apache ActiveMQ Artemis messaging server integrated in a remote instance of JBoss EAP. The size of the client thread pool is set using the **activemq.artemis.client.global.thread.pool.max.size** system property.

However, the main use of the client thread pool is on the client side. A client can configure each **ClientSessionFactory** instance to not use the provided pool, but instead to create and maintain its own client thread pool. Any sessions created from that **ClientSessionFactory** will use the newly created pool.

To configure a **ClientSessionFactory** instance to use its own pools, use the appropriate setter methods immediately after creation. For example:

```
ServerLocator locator = ActiveMQClient.createServerLocatorWithoutHA(...)
ClientSessionFactory myFactory = locator.createClientSessionFactory();
myFactory.setUseGlobalPools(false);
myFactory.setScheduledThreadPoolMaxSize(10);
myFactory.setThreadPoolMaxSize(-1);
```

If you are using the JMS API, you can set the same parameters on the **ClientSessionFactory**. For example:

```
ConnectionFactory myFactory =
ActiveMQJMSClient.createConnectionFactory(myFactory);
myFactory.isUseGlobalPools(false);
myFactory.setScheduledThreadPoolMaxSize(10);
myFactory.setThreadPoolMaxSize(-1);
```

If you are using JNDI to instantiate **ActiveMQConnectionFactory** instances, you can also set these parameters using the management CLI, as in the examples given below for a standalone instance of JBoss EAP.

```
/subsystem=messaging-activemq/server=default/connection-
factory=myConnectionFactory:write-attribute(name=use-global-
pools,value=false)

/subsystem=messaging-activemq/server=default/connection-
```

```
factory=myConnectionFactory:write-attribute(name=scheduled-thread-pool-  
max-size,value=10)
```

```
/subsystem=messaging-activemq/server=default/connection-  
factory=myConnectionFactory:write-attribute(name=thread-pool-max-  
size,value=1)
```

Note that the management CLI will remind you that a reload of the instance is required after you execute each of the above commands.

## CHAPTER 25. CONFIGURING DUPLICATE MESSAGE DETECTION

When a sender sends a message to another server, there can be a situation where the target server or the connection fails after sending the message, but before sending a response to the sender indicating that the process was successful. In these situations, it is very difficult for the sender to determine whether the message was sent successfully to the intended receiver. If the sender decides to resend the last message, it can result in a duplicate message being sent to the address.

You can configure duplicate message detection in JBoss EAP messaging so that your application does not need to provide the logic to filter duplicate messages.

### 25.1. USING DUPLICATE MESSAGE DETECTION FOR SENDING MESSAGES

To enable duplicate message detection for sent messages, you need to set the value of the `org.apache.activemq.artemis.api.core.Message.HDR_DUPLICATE_DETECTION_ID` property, which resolves to `_AMQ_DUPL_ID`, to a unique value. When the target server receives the messages, if the `_AMQ_DUPL_ID` property is set, it will check its memory cache to see if it has already received a message with the value of that header. If it has, then this message will be ignored. See [Configuring the Duplicate ID Cache](#) for more information.

The value of the `_AMQ_DUPL_ID` property can be of type `byte[]` or `SimpleString` if you are using the core API. If you are using JMS, it must be a `String`.

The following example shows how to set the property for core API.


```
SimpleString myUniqueID = "This is my unique id"; // Can use a UUID for
this

ClientMessage message = session.createMessage(true);
message.setStringProperty(HDR_DUPLICATE_DETECTION_ID, myUniqueID);
```

The following example shows how to set the property for JMS clients.

```
String myUniqueID = "This is my unique id"; // Can use a UUID for this

Message jmsMessage = session.createMessage();
message.setStringProperty(HDR_DUPLICATE_DETECTION_ID.toString(),
myUniqueID);
```



#### IMPORTANT

Duplicate messages are *not* detected when they are sent within the same transaction using the `HDR_DUPLICATE_DETECTION_ID` property.

### 25.2. CONFIGURING THE DUPLICATE ID CACHE

The server maintains caches of received values of the `_AMQ_DUPL_ID` property that is sent to each address. Each address maintains its own address cache.

The cache is fixed in terms of size. The maximum size of cache is configured using the `id-cache-`

**size** attribute. The default value of this parameter is **20000** elements. If the cache has a maximum size of **n** elements, then the **(n + 1)**th ID stored will overwrite the element **0** in the cache. The value is set using the following management CLI command:

```
/subsystem=messaging-activemq/server=default:write-attribute(name=id-cache-size,value=SIZE)
```

The caches can also be configured to persist to disk. This can be configured by setting the **persist-id-cache** attribute using the following management CLI command.

```
/subsystem=messaging-activemq/server=default:write-attribute(name=persist-id-cache,value=true)
```

If this value is set to **true**, then each ID will be persisted to permanent storage as they are received. The default value for this parameter is **true**.



#### NOTE

Set the size of the duplicate ID cache to a large size in order to ensure that resending of messages does not overwrite the previously sent messages stored in the cache.

## CHAPTER 26. HANDLING SLOW CONSUMERS

A slow consumer with a server-side queue can pose a significant problem for server performance. As messages build up in the consumer's server-side queue, memory usage will increase. Consequently, the server may enter paging mode which can impact performance. Criteria can be set, however, so that consumers that do not acknowledge messages quickly enough can be disconnected from the server. In the case of a non-durable JMS subscriber, this means the server would remove the subscription and all of its messages, freeing up valuable server resources.

The calculation to determine whether or not a consumer is slow inspects only the number of messages that a particular consumer has acknowledged. It does not take into account whether flow control has been enabled on the consumer, or whether the consumer is streaming a large message, for example. Keep this in mind when configuring slow consumer detection.

Slow consumer checks are performed using the scheduled thread pool. Each queue on the server with slow consumer detection enabled will cause a new entry in the internal `java.util.concurrent.ScheduledThreadPoolExecutor` instance. If there are a high number of queues and the `slow-consumer-check-period` is relatively low, then there may be delays in executing some of the checks. However, this will not impact the accuracy of the calculations used by the detection algorithm. See [Thread Management](#) for more details about this thread pool.

Slow consumer handling is on a per `address-setting` basis. See [Address Settings](#) for more information on configuring an `address-setting`, and refer to the appendix for the list of `address-setting` attributes. There are three attributes used to configure the handling of slow consumers. They are:

### `slow-consumer-check-period`

How often to check, in seconds, for slow consumers. The default is `5`.

### `slow-consumer-policy`

Determines what happens when a slow consumer is identified. The valid options are `KILL` or `NOTIFY`:

- `KILL` will kill the consumer's connection, which will impact any client threads using that same connection.
- `NOTIFY` will send a `CONSUMER_SLOW` management notification to the client.

The default is `NOTIFY`.

### `slow-consumer-threshold`

The minimum rate of message consumption allowed before a consumer is considered slow. The default is `-1`, which is unbounded.

Use the management CLI to read the current values for any of the attributes. For example, use the following command to read the current `slow-consumer-policy` for the `address-setting` with the name `myAddress`.

```
/subsystem=messaging-activemq/server=default/address-setting=myAddress:read-attribute(name=slow-consumer-policy)
```

Likewise, use the following example to set the same `slow-consumer-policy`.

```
/subsystem=messaging-activemq/server=default/address-  
setting=myAddress:write-attribute(name=slow-consumer-policy,value=  
<NEW_VALUE>)
```



## **PART III. CONFIGURING MULTI-NODE MESSAGING SYSTEMS**

## CHAPTER 27. CONFIGURING JMS BRIDGES

JBoss EAP messaging includes a fully functional JMS message bridge. The function of a JMS bridge is to consume messages from a source queue or topic and send them to a target queue or topic, typically on a different server.

If the source destination is a topic, the JMS bridge creates a subscription for it. If the `client-id` and `subscription-name` attributes are configured for the JMS bridge, the subscription is durable. This means that no messages are missed if the JMS bridge is stopped and then restarted.

The source and target servers do not have to be in the same cluster, which makes bridging suitable for reliably sending messages from one cluster to another, for instance across a WAN, and where the connection may be unreliable.



### NOTE

Do not confuse a JMS bridge with a core bridge. A JMS bridge can be used to bridge any two JMS 1.1 compliant JMS providers and uses the JMS API. A [core bridge](#) is used to bridge any two JBoss EAP messaging instances and uses the core API. It is preferable to use a core bridge instead of a JMS bridge whenever possible.

Below is an example configuration of a JBoss EAP JMS bridge.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
  </server>
  <jms-bridge name="my-jms-bridge" max-batch-time="100" max-batch-size="10" max-retries="1" failure-retry-interval="500" quality-of-service="AT_MOST_ONCE">
    <source destination="jms/queue/InQueue" connection-factory="ConnectionFactory">
      <source-context/>
    </source>
    <target destination="jms/queue/OutQueue" connection-factory="jms/RemoteConnectionFactory">
      <target-context>
        <property name="java.naming.factory.initial" value="org.jboss.naming.remote.client.InitialContextFactory"/>
        <property name="java.naming.provider.url" value="http-remoting://192.168.40.1:8080"/>
      </target-context>
    </target>
  </jms-bridge>
  ...
</subsystem>
```

### Adding a JMS Bridge Using the Management CLI

A JMS bridge can be added using the following management CLI command. Note that the source and target destinations must already be defined in the configuration. See the [table in the appendix](#) for a full list of configurable attributes.

```
/subsystem=messaging-activemq/jms-bridge=my-jms-bridge:add(quality-of-service=AT_MOST_ONCE,failure-retry-interval=500,max-retries=1,max-batch-
```

```
size=10,max-batch-time=100,source-connection-
factory=ConnectionFactory,source-destination=jms/queue/InQueue,source-
context={},target-connection-factory=jms/RemoteConnectionFactory,target-
destination=jms/queue/OutQueue,target-context=
{java.naming.factory.initial=org.jboss.naming.remote.client.InitialContext
Factory,java.naming.provider.url=http-remoting://192.168.40.1:8080})
```



## NOTE

When shutting down a server that has a deployed JMS bridge with the **quality-of-service** attribute set to **ONCE\_AND\_ONLY\_ONCE**, be sure to shut the server down with the JMS bridge first to avoid unexpected errors.

## 27.1. QUALITY OF SERVICE

In JBoss EAP, **quality-of-service** is a configurable attribute that determines how messages are consumed and acknowledged. The valid values for **quality-of-service** and their descriptions are below. See the [table in the appendix](#) for a full list of JMS bridge attributes.

### AT\_MOST\_ONCE

Messages will reach the destination from the source at the most one time. The messages are consumed from the source and acknowledged before sending to the destination. Therefore, there is a possibility that messages could be lost if a failure occurs between their removal from the source and their arrival at the destination. This mode is the default value.

This mode is available for both durable and non-durable messages.

### DUPLICATES\_OK

Messages are consumed from the source and then acknowledged after they have been successfully sent to the destination. Therefore, there is a possibility that messages could be sent again if a failure occurs after the initial message was sent but before it is acknowledged.

This mode is available for both durable and non-durable messages.

### ONCE\_AND\_ONLY\_ONCE

Messages will reach the destination from the source once and only once. If both the source and the destination are on the same server instance, this can be achieved by sending and acknowledging the messages in the same local transaction. If the source and destination are on different servers, this is achieved by enlisting the sending and consuming sessions in a JTA transaction. The JTA transaction is controlled by a JTA Transaction Manager which will need to be set using the **setTransactionManager()** method on the bridge.

This mode is only available for durable messages.



## WARNING

When shutting down a server that has a deployed JMS bridge with the **quality-of-service** attribute set to **ONCE\_AND\_ONLY\_ONCE**, be sure to shut the server down with the JMS bridge first to avoid unexpected errors.

It may be possible to provide once and only once semantics by using the **DUPLICATES\_OK** mode instead of **ONCE\_AND\_ONLY\_ONCE** and then checking for duplicates at the destination and discarding them. See [Configuring Duplicate Message Detection](#) for more information. However, the cache would only be valid for a certain period of time. This approach therefore is not as watertight as using **ONCE\_AND\_ONLY\_ONCE**, but it may be a good choice depending on your specific application needs.

## 27.2. TIMEOUTS AND THE JMS BRIDGE

There is a possibility that the target or source server will not be available at some point in time. If this occurs, the bridge will try to reconnect a number of times equal to the value of **max-retries**. The wait between attempts is set by **failure-retry-interval**.



### WARNING

Because each JMS bridge has its own **max-retries** parameter, you should use a connection factory that does not set the **reconnect-attempts** parameter, or sets it to **0**. This will avoid a potential collision that may result in longer reconnection times. Also note that any connection factory referenced by a JMS bridge with the **quality-of-service** set to **ONCE\_AND\_ONLY\_ONCE** needs to have the **factory-type** set to **XA\_GENERIC**, **XA\_TOPIC**, or **XA\_QUEUE**.

However, it is possible for the JNDI lookup to hang if the network were to disappear during the JNDI lookup. To stop this from occurring, configure the JNDI definition to time out by setting the following properties on the JNDI **InitialContext**.

#### **jnp.timeout**

Sets the connection timeout for the initial connection

#### **jnp.sotimeout**

Sets the read timeout for the socket.

Once the initial JNDI connection has succeeded all subsequent calls are made using RMI. If you want to control the timeouts for the RMI connections, then this can be done with system properties. JBoss EAP uses Sun's RMI, and the properties can be found in the [Java SE documentation](#). The default connection timeout is 10 seconds and the default read timeout is 18 seconds.

If you implement your own factories for looking up JMS resources, you will have to manage any timeout issues in the implementation and expose configuration properties accordingly.

## CHAPTER 28. CONFIGURING CORE BRIDGES

The function of a bridge is to consume messages from one destination and forward them to another one, typically on a different JBoss EAP messaging server.

The source and target servers do not have to be in the same cluster which makes bridging suitable for reliably sending messages from one cluster to another, for instance across a WAN, or internet and where the connection may be unreliable.

The bridge has built-in resilience to failure so if the target server connection is lost, for example, due to network failure, the bridge will retry connecting to the target until it comes back online. When it comes back online it will resume operation as normal.

Bridges are a way to reliably connect two separate JBoss EAP messaging servers together. With a core bridge both source and target servers must be JBoss EAP 7 messaging servers.



### NOTE

Do not confuse a core bridge with a JMS bridge. A core bridge is used to bridge any two JBoss EAP messaging instances and uses the core API. A [JMS bridge](#) can be used to bridge any two JMS 1.1 compliant JMS providers and uses the JMS API. It is preferable to use a core bridge instead of a JMS bridge whenever possible.

Below is an example configuration of a JBoss EAP messaging core bridge.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <bridge name="my-core-bridge" static-connectors="bridge-connector"
queue-name="jms.queue.InQueue"/>
    ...
  </server>
</subsystem>
```

This core bridge can be added using the following management CLI command. Note that when defining a core bridge, you must define a **queue-name** and either **static-connectors** or **discovery-group**. See the [table in the appendix](#) for a full list of configurable attributes.

```
/subsystem=messaging-activemq/server=default/bridge=my-core-
bridge:add(static-connectors=[bridge-connector], queue-
name=jms.queue.InQueue)
```

### 28.1. CONFIGURING A CORE BRIDGE FOR DUPLICATE DETECTION

Core bridges can be configured to automatically add a unique duplicate ID value, if there is not already one in the message, before forwarding the message to the target. To configure a core bridge for duplicate message detection set the **use-duplicate-detection** attribute to **true**, which is the default value.

```
/subsystem=messaging-activemq/server=default/bridge=my-core-bridge:write-
attribute(name=use-duplicate-detection,value=true)
```

## CHAPTER 29. CLUSTERS OVERVIEW

JBoss EAP messaging clusters allow groups of JBoss EAP messaging servers to be grouped together in order to share message processing load. Each active node in the cluster is an active JBoss EAP messaging server which manages its own messages and handles its own connections.

The cluster is formed by each node declaring cluster connections to other nodes in the JBoss EAP configuration file. When a node forms a cluster connection to another node, it internally creates a [core bridge](#) connection between itself and the other node. This is done transparently behind the scenes; you do not have to declare an explicit bridge for each node. These cluster connections allow messages to flow between the nodes of the cluster to balance the load.

An important part of clustering is [server discovery](#) where servers can broadcast their connection details so clients or other servers can connect to them with minimum configuration.

This section also discusses [client-side load balancing](#), to balance client connections across the nodes of the cluster, and [message redistribution](#), where JBoss EAP messaging will redistribute messages between nodes to avoid starvation.



### WARNING

Once a cluster node has been configured, it is common to simply copy that configuration to other nodes to produce a symmetric cluster.

In fact, each node in the cluster must share the same configuration for the following elements in order to avoid unexpected errors:

- cluster-connection
- broadcast-group
- discovery-group
- address-settings, including queues and topics

However, care must be taken when copying the JBoss EAP messaging files. Do not copy the messaging data, the bindings, journal, and large-messages directories from one node to another. When a cluster node is started for the first time and initializes its journal files, it persists a special identifier to the journal directory. The identifier must be unique among nodes for the cluster to form properly.

### 29.1. SERVER DISCOVERY

Server discovery is a mechanism by which servers can propagate their connection details to:

- Messaging clients  
A messaging client wants to be able to connect to the servers of the cluster without having specific knowledge of which servers in the cluster are up at any one time.
- Other servers.

Servers in a cluster want to be able to create cluster connections to each other without having prior knowledge of all the other servers in the cluster.

This information, or cluster topology, is sent around normal JBoss EAP messaging connections to clients and to other servers over cluster connections. However, you need a way to establish the initial first connection. This can be done using dynamic discovery techniques like UDP and JGroups, or by providing a [static list](#) of initial connectors.

### 29.1.1. Broadcast Groups

A broadcast group is the means by which a server broadcasts connectors over the network. A [connector](#) defines a way in which a client, or other server, can make connections to the server.

The broadcast group takes a set of connectors and broadcasts them on the network. Depending on which broadcasting technique you configure the cluster, it uses either UDP or JGroups to broadcast connector pairs information.

Broadcast groups are defined in the **messaging-activemq** subsystem of the server configuration. There can be many broadcast groups per JBoss EAP messaging server.

#### Configure a Broadcast Group Using UDP

Below is an example configuration of a messaging server that defines a UDP broadcast group. Note that this configuration relies on a **messaging-group** socket binding.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <broadcast-group name="my-broadcast-group" connectors="http-connector"
socket-binding="messaging-group"/>
    ...
  </server>
</subsystem>
...
<socket-binding-group name="standard-sockets" default-interface="public"
port-offset="{jboss.socket.binding.port-offset:0}">
  ...
  <socket-binding name="messaging-group" interface="private" port="5432"
multicast-address="231.7.7.7" multicast-port="9876"/>
  ...
</socket-binding-group>
```

This configuration can be achieved using the following management CLI commands:

1. Add the **messaging-group** socket binding.

```
/socket-binding-group=standard-sockets/socket-binding=messaging-
group:add(interface=private,port=5432,multicast-
address=231.7.7.7,multicast-port=9876)
```

2. Add the broadcast group.

```
/subsystem=messaging-activemq/server=default/broadcast-group=my-
broadcast-group:add(socket-binding=messaging-group,broadcast-
period=2000,connectors=[http-connector])
```

## Configure a Broadcast Group Using JGroups

Below is an example configuration of a messaging server that defines broadcast group that uses the default JGroups broadcast group, which uses UDP. Note that to be able to use JGroups to broadcast, you must set a `jgroups-channel`.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <broadcast-group name="my-broadcast-group" connectors="http-connector"
jgroups-channel="activemq-cluster"/>
    ...
  </server>
</subsystem>
```

This can be configured using the following management CLI command:

```
/subsystem=messaging-activemq/server=default/broadcast-group=my-broadcast-
group:add(connectors=[http-connector],jgroups-channel=activemq-cluster)
```

## Broadcast Group Attributes

The below table lists the configurable attributes for a broadcast group.

Attribute	Description
broadcast-period	The period in milliseconds between consecutive broadcasts.
connectors	The names of connectors that will be broadcast.
jgroups-channel	The name used by a JGroups channel to join a cluster.
jgroups-stack	The name of a stack defined in the <b>jgroups</b> subsystem that is used to form a cluster.
socket-binding	The broadcast group socket binding.

### 29.1.2. Discovery Groups

While the broadcast group defines how connector information is broadcasted from a server, a discovery group defines how connector information is received from a broadcast endpoint, for example, a UDP multicast address or JGroup channel.

A discovery group maintains a list of connectors, one for each broadcast by a different server. As it receives broadcasts on the broadcast endpoint from a particular server, it updates its entry in the list for that server. If it has not received a broadcast from a particular server for a length of time it will remove that server's entry from its list.

Discovery groups are used in two places in JBoss EAP messaging:

- By cluster connections so they know how to obtain an initial connection to download the topology.
- By messaging clients so they know how to obtain an initial connection to download the topology.



Although a discovery group will always accept broadcasts, its current list of available live and backup servers is only ever used when an initial connection is made. From then on, server discovery is done over the normal JBoss EAP messaging connections.



## NOTE

Each discovery group must be configured with a broadcast endpoint (UDP or JGroups) that matches its broadcast group counterpart. For example, if the broadcast group is configured using UDP, the discovery group must also use UDP and the same multicast address.

### 29.1.2.1. Configure Discovery Groups on the Server

Discovery groups are defined in the **messaging-activemq** subsystem of the server configuration. There can be many discovery groups per JBoss EAP messaging server.

#### Configure a Discovery Group Using UDP

Below is an example configuration of a messaging server that defines a UDP discovery group. Note that this configuration relies on a **messaging-group** socket binding.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <discovery-group name="my-discovery-group" refresh-timeout="10000"
socket-binding="messaging-group"/>
    ...
  </server>
</subsystem>
...
<socket-binding-group name="standard-sockets" default-interface="public"
port-offset="{jboss.socket.binding.port-offset:0}">
  ...
  <socket-binding name="messaging-group" interface="private" port="5432"
multicast-address="231.7.7.7" multicast-port="9876"/>
  ...
</socket-binding-group>
```

This configuration can be achieved using the following management CLI commands:

1. Add the **messaging-group** socket binding.

```
/socket-binding-group=standard-sockets/socket-binding=messaging-
group:add(interface=private,port=5432,multicast-
address=231.7.7.7,multicast-port=9876)
```

2. Add the discovery group.

```
/subsystem=messaging-activemq/server=default/discovery-group=my-
discovery-group:add(socket-binding=messaging-group,refresh-
timeout=10000)
```

#### Configure a Discovery Group Using JGroups

Below is an example configuration of a messaging server that defines a JGroups discovery group.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <discovery-group name="my-discovery-group" refresh-timeout="10000"
jgroups-channel="activemq-cluster"/>
    ...
  </server>
</subsystem>
```

This can be configured using the following management CLI command:

```
/subsystem=messaging-activemq/server=default/discovery-group=my-discovery-
group:add(refresh-timeout=10000,jgroups-channel=activemq-cluster)
```

### Discovery Group Attributes

The below table lists the configurable attributes for a discovery group.

Attribute	Description
initial-wait-timeout	Period, in milliseconds, to wait for an initial broadcast to give us at least one node in the cluster.
jgroups-channel	The name used by a JGroups channel to join a cluster.
jgroups-stack	The name of a stack defined in the <b>jgroups</b> subsystem that is used to form a cluster.
refresh-timeout	Period the discovery group waits after receiving the last broadcast from a particular server before removing that server's connector pair entry from its list.
socket-binding	The discovery group socket binding.



#### WARNING

The JGroups attributes and UDP-specific attributes described above are exclusive of each other. Only one set can be specified in a discovery group configuration.

### 29.1.2.2. Configure Discovery Groups on the Client Side

You can use JMS or the core API to configure a JBoss EAP messaging client to discover a list of servers to which it can connect.

#### Configure Client Discovery using JMS

Clients using JMS can look up the relevant **ConnectionFactory** with JNDI. The **entries** attribute of a **connection-factory** or a **pooled-connection-factory** specifies the JNDI name under which the factory will be exposed. Below is an example of a **ConnectionFactory** configured for a remote

client to lookup with JNDI:

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <connection-factory name="RemoteConnectionFactory"
entries="java:jboss/exported/jms/RemoteConnectionFactory"
connectors="http-connector"/>
    ...
  </server>
</subsystem>
```

## NOTE

It is important to remember that only JNDI names bound in the **java:jboss/exported** namespace are available to remote clients. If a **connection-factory** has an entry bound in the **java:jboss/exported** namespace a remote client would look up the **connection-factory** using the text after **java:jboss/exported**. For example, the **RemoteConnectionFactory** is bound by default to **java:jboss/exported/jms/RemoteConnectionFactory** which means a remote client would look-up this **connection-factory** using **jms/RemoteConnectionFactory**. A **pooled-connection-factory** should not have any entry bound in the **java:jboss/exported** namespace because a **pooled-connection-factory** is not suitable for remote clients.

Since JMS 2.0, a default JMS connection factory is accessible to Java EE applications under the JNDI name **java:comp/DefaultJMSConnectionFactory**. The JBoss EAP **messaging-activemq** subsystem defines a **pooled-connection-factory** that is used to provide this default connection factory. Any parameter change on this **pooled-connection-factory** will be taken into account by any Java EE application looking the default JMS provider under the JNDI name **java:comp/DefaultJMSConnectionFactory**. Below is the default pooled connection factory as defined in the **\*-full** and **\*-full-ha** profiles.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <pooled-connection-factory name="activemq-ra" transaction="xa"
entries="java:/JmsXA java:jboss/DefaultJMSConnectionFactory"
connectors="in-vm"/>
    ...
  </server>
</subsystem>
```

## Configure Client Discovery using the Core API

If you are using the core API to directly instantiate **ClientSessionFactory** instances, then you can specify the discovery group parameters directly when creating the session factory. For example:

```
final String groupAddress = "231.7.7.7";
final int groupPort = 9876;

ServerLocator factory =
  ActiveMQClient.createServerLocatorWithHA(new
DiscoveryGroupConfiguration(
```

```

        groupAddress,
        groupPort,
        new UDPBroadcastGroupConfiguration(groupAddress, groupPort, null,
-1)));
ClientSessionFactory factory = locator.createSessionFactory();
ClientSession session1 = factory.createSession();
ClientSession session2 = factory.createSession();

```

You can use the `setDiscoveryRefreshTimeout()` setter method on the `DiscoveryGroupConfiguration` to set the `refresh-timeout` value, which defaults to **10000** milliseconds.

You can also use the `setDiscoveryInitialWaitTimeout()` setter method on the `DiscoveryGroupConfiguration` to set the `initial-wait-timeout` value, which determines how long the session factory will wait before creating the first session. The default value is **10000** milliseconds.

### 29.1.3. Static Discovery

In situations where you can not or do not want to use UDP on your network, you can configure a connection with an initial list of one or more servers.

This does not mean that you have to know where all your servers are going to be hosted. You can configure these servers to connect to a reliable server, and have their connection details propagated by way of that server.

#### Configuring a Cluster Connection

For cluster connections there, is no additional configuration needed, you just need to make sure that any [connectors](#) are defined in the usual manner. These are then referenced by the cluster connection configuration.

#### Configuring a Client Connection

A static list of possible servers can also be used by a client.

#### Configuring Client Discovery Using JMS

The recommended way to use static discovery with JMS is to configure a `connection-factory` with multiple connectors (each pointing to a unique node in the cluster) and have the client look up the `ConnectionFactory` using JNDI. Below is a snippet of configuration showing just such a `connection-factory`:

```

<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <connection-factory name="MyConnectionFactory" entries="..."
connectors="http-connector http-node1 http-node2"/>
    ...
  </server>
</subsystem>

```

In the above example, `http-connector` is an HTTP connector (`<http-connector>`) pointing to the local server, `http-node1` is an HTTP connector pointing to server `node1`, and so on. See the [Connectors and Acceptors](#) section for configuring connectors in the server configuration.

#### Configuring Client Discovery Using the Core API

If you are using the core API, create a unique **TransportConfiguration** for each server in the cluster and pass them into the method responsible for creating the **ServerLocator**, as in the below example code.

```
HashMap<String, Object> map = new HashMap<String, Object>();
map.put("host", "myhost");
map.put("port", "8080");

HashMap<String, Object> map2 = new HashMap<String, Object>();
map2.put("host", "myhost2");
map2.put("port", "8080");

TransportConfiguration server1 = new
TransportConfiguration(NettyConnectorFactory.class.getName(), map);
TransportConfiguration server2 = new
TransportConfiguration(NettyConnectorFactory.class.getName(), map2);

ServerLocator locator = ActiveMQClient.createServerLocatorWithHA(server1,
server2);
ClientSessionFactory factory = locator.createSessionFactory();
ClientSession session = factory.createSession();
```

## 29.2. SERVER-SIDE MESSAGE LOAD BALANCING

If a cluster connection is defined between nodes of a cluster, then JBoss EAP messaging will load balance messages arriving at a particular node from a client.

A messaging cluster connection can be configured to load balance messages in a round robin fashion, irrespective of whether there are any matching consumers on other nodes. It can also be configured to distribute to other nodes only if matching consumers exist. See the [Message Redistribution](#) section for more information.

### Configuring the Cluster Connection

A cluster connection group servers into clusters so that messages can be load balanced between the nodes of the cluster. A cluster connection is defined in the JBoss EAP server configuration using the **cluster-connection** element.



#### WARNING

Red Hat supports using only one **cluster-connection** within the **messaging-activemq** subsystem.

Below is the default **cluster-connection** as defined in the **\*-full** and **\*-full-ha** profiles. See [Cluster Connection Attributes](#) for the complete list of attributes.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <cluster-connection name="my-cluster" discovery-group="dg-group1"
```

```
connector-name="http-connector" address="jms"/>
    ...
</server>
</subsystem>
```

In the case shown above the cluster connection will load balance messages sent to addresses that start with "jms". This cluster connection will, in effect, apply to all JMS queues and topics since they map to core queues that start with the substring "jms".

## NOTE

When a packet is sent using a cluster connection and is at a blocked state and waiting for acknowledgements, the **call-timeout** attribute specifies how long it will wait for the reply before throwing an exception. The default value is **30000**. In certain cases, for example, if the remote JMS broker is disconnected from network and the transaction is incomplete, the thread could remain stuck until connection is re-established. To avoid this situation, it is recommended to use the **call-failover-timeout** attribute along with the **call-timeout** attribute. The **call-failover-timeout** attribute is used when a call is made during a failover attempt. The default value is **-1**, which means no timeout. For more information on Client Failover, see [Automatic Client Failover](#).

## NOTE

Alternatively, if you would like the cluster connection to use a static list of servers for discovery then you can use the **static-connectors** attribute. For example:

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <cluster-connection name="my-cluster" static-
connectors="server0-connector server1-connector" .../>
    ...
  </server>
</subsystem>
```

In this example, there are two servers defined where we know that at least one will be available. There may be many more servers in the cluster, but these will be discovered using one of these connectors once an initial connection has been made.

## Configuring a Cluster Connection for Duplicate Detection

The cluster connection internally uses a [core bridge](#) to move messages between nodes of the cluster. To configure a cluster connection for duplicate message detection, set the **use-duplicate-detection** attribute to **true**, which is the default value.

```
/subsystem=messaging-activemq/server=default/cluster-connection=my-
cluster:write-attribute(name=use-duplicate-detection,value=true)
```

## Cluster User Credentials

When creating connections between nodes of a cluster to form a cluster connection, JBoss EAP messaging uses a cluster user and password.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
```

```

...
<cluster password="{jboss.messaging.cluster.password:CHANGE ME!!}" />
...
</server>
</subsystem>

```



### WARNING

The default value for `cluster-user` is `ACTIVEMQ.CLUSTER.ADMIN.USER` and the default value for `cluster-password` is `CHANGE ME!!`. It is imperative that these values are changed from their default, or remote clients will be able to make connections to the server using the default values. If they are not changed from the default, JBoss EAP messaging will detect this and display a warning upon every startup.

## 29.3. CLIENT-SIDE LOAD BALANCING

With JBoss EAP messaging client-side load balancing, subsequent sessions created using a single session factory can be connected to different nodes of the cluster. This allows sessions to spread smoothly across the nodes of a cluster and not be clumped on any particular node.

The recommended way to declare a load balancing policy to be used by the client factory is to set the `connection-load-balancing-policy-class-name` attribute of the `<connection-factory>` resource. JBoss EAP messaging provides the following out-of-the-box load balancing policies, and you can also implement your own.

### Round robin

With this policy, the first node is chosen randomly then each subsequent node is chosen sequentially in the same order.

For example, nodes might be chosen in the order `B, C, D, A, B, C, D, A, B` or `D, A, B, C, D, A, B, C`.

Use

`org.apache.activemq.artemis.api.core.client.loadbalance.RoundRobinConnectionLoadBalancingPolicy` as the `connection-load-balancing-policy-class-name`.

### Random

With this policy, each node is chosen randomly.

Use

`org.apache.activemq.artemis.api.core.client.loadbalance.RandomConnectionLoadBalancingPolicy` as the `connection-load-balancing-policy-class-name`.

### Random Sticky

With this policy, the first node is chosen randomly and then reused for subsequent connections.

Use

`org.apache.activemq.artemis.api.core.client.loadbalance.RandomStickyConnectionLoadBalancingPolicy` as the `connection-load-balancing-policy-class-name`.

### First Element



With this policy, the first, or 0th, node is always returned.

Use

```
org.apache.activemq.artemis.api.core.client.loadbalance.FirstElementConnectionLoadBalancingPolicy as the connection-load-balancing-policy-class-name .
```

You can also implement your own policy by implementing the interface **org.apache.activemq.artemis.api.core.client.loadbalance.ConnectionLoadBalancingPolicy**

## 29.4. MESSAGE REDISTRIBUTION

With message redistribution, JBoss EAP messaging can be configured to automatically redistribute messages from queues which have no consumers back to other nodes in the cluster which do have matching consumers. To enable this functionality, cluster connection's **message-load-balancing-type** must be set to **ON\_DEMAND**, which is the default value. You can set this using the following management CLI command.

```
/subsystem=messaging-activemq/server=default/cluster-connection=my-cluster:write-attribute(name=message-load-balancing-type,value=ON_DEMAND)
```

Message redistribution can be configured to kick in immediately after the last consumer on a queue is closed, or to wait a configurable delay after the last consumer on a queue is closed before redistributing. This is configured using the **redistribution-delay** attribute.

You use the **redistribution-delay** attribute to set how many milliseconds to wait after the last consumer is closed on a queue before redistributing messages from that queue to other nodes of the cluster that have matching consumers. A value of **-1**, which is the default value, means that messages will never be redistributed. A value of **0** means that messages will be immediately redistributed.

The **address-setting** in the default JBoss EAP configuration sets a **redistribution-delay** value of **1000**, meaning that it will wait 1000 milliseconds before redistributing messages.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
    ...
    <address-setting name="#" redistribution-delay="1000" message-counter-history-day-limit="10" page-size-bytes="2097152" max-size-bytes="10485760" expiry-address="jms.queue.ExpiryQueue" dead-letter-address="jms.queue.DLQ"/>
    ...
  </server>
</subsystem>
```

It often makes sense to introduce a delay before redistributing as it is a common case that a consumer closes but another one quickly is created on the same queue. In this case, you may not want to redistribute immediately since the new consumer will arrive shortly.

Below is an example of an **address-setting** that sets a **redistribution-delay** of **0** for any queue or topic that is bound to an address that starts with "jms.". In this case, messages will be redistributed immediately.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
  <server name="default">
```



```

...
<address-setting name="jms.#" redistribution-delay="0"/>
...
</server>
</subsystem>

```


This address setting can be added using the following management CLI command.

```

/subsystem=messaging-activemq/server=default/address-
setting=jms.#:add(redistribution-delay=1000)

```

## 29.5. CLUSTERED MESSAGE GROUPING



### IMPORTANT

This feature is provided as Technology Preview only. It is not supported for use in a production environment, and it may be subject to significant future changes. See [Technology Preview Features Support Scope](#) on the Red Hat Customer Portal for information about the support scope for Technology Preview features.

Clustered grouping follows a different approach relative to normal [message grouping](#). In a cluster, message groups with specific group ids can arrive on any of the nodes. It is important for a node to determine which group ids are bound to which consumer on which node. Each node is responsible for routing message groups correctly to the node which has the consumer processing those group ids, irrespective of where the message groups arrive by default. Once messages with a given group id are sent to a specific consumer connected to the given node in the cluster, then those messages are never sent to another node even if the consumer is disconnected.

This situation is addressed by a grouping handler. Each node has a grouping handler and this grouping handler (along with other handlers) is responsible for routing the message groups to the correct node. There are two types of grouping handlers: **LOCAL** and **REMOTE**.

The local handler is responsible for deciding the route that a message group should take. The remote handlers communicate with the local handler and work accordingly. Each cluster should choose a specific node to have a local grouping handler and all the other nodes should have remote handlers.



### WARNING

If message grouping is used in a cluster, it will break if a node configured as a remote grouping handler fails. Setting up a backup for the remote grouping handler will not correct this.

The node that initially receives a message group takes the routing decision based on regular cluster routing conditions (round-robin queue availability). The node proposes this decision to the respective grouping handler which then routes the messages to the proposed queue if it accepts the proposal.

If the grouping handler rejects the proposal, it proposes some other route and the routing takes place accordingly. The other nodes follow suite and forward the message groups to the chosen queue. After a message arrives on a queue, it is pinned to a customer on that queue.

You can configure grouping handlers using the management CLI. The following command adds a **LOCAL** grouping handler with the address `news.europe.#`.

```
/subsystem=messaging-activemq/server=default/grouping-handler=my-group-handler:add(grouping-handler-address="news.europe.#", type=LOCAL)
```

This will require a server reload.

```
reload
```

The below table lists the configurable attributes for a **grouping-handler**.

Attribute	Description
group-timeout	With a <b>REMOTE</b> handler, this value specifies how often the <b>REMOTE</b> will notify the <b>LOCAL</b> that the route was used. With a <b>LOCAL</b> handler, if a route is not used for the time specified, it is removed, and a new path would need to be established. The value is in milliseconds.
grouping-handler-address	A reference to a cluster connection and the address it uses.
reaper-period	How often the reaper will be run to check for timed out group bindings (only valid for <b>LOCAL</b> handlers).
timeout	How long to wait for a handling decision to be made; an exception will be thrown during the send if this timeout is reached, ensuring that strict ordering is kept.
type	Whether the handler is the single local handler for the cluster, which makes handling decisions, or a remote handler which converses with the local handler. Possible values are <b>LOCAL</b> or <b>REMOTE</b> .

### 29.5.1. Best Practices for Clustered Message Grouping


Some best practices for clustered grouping are as follows:

- If you create and close consumers regularly, make sure that your consumers are distributed evenly across the different nodes. Once a queue is pinned, messages are automatically transferred to that queue regardless of removing customers from it.
- If you wish to remove a queue that has a message group bound to it, make sure the queue is deleted by the session that is sending the messages. Doing this will ensure that other nodes will not try to route messages to this queue after it is removed.
- As a failover mechanism, always replicate the node that has the local grouping handler.

## CHAPTER 30. HIGH AVAILABILITY

High availability is the ability for the system to continue functioning after failure of one or more of the servers.

A part of high availability is failover which is the ability for client connections to migrate from one server to another in event of server failure so client applications can continue to operate.




### NOTE

Only persistent message data will survive failover. Any non persistent message data will not be available after failover.

### 30.1. LIVE / BACKUP PAIRS

JBoss EAP 7 messaging allows servers to be linked together as live - backup pairs where each live server has a backup. Live servers receive messages from clients, while a backup server is not operational until failover occurs. A backup server can be owned by only one live server, and it will remain in passive mode, waiting to take over the live server's work.



### NOTE

There is a one-to-one relation between a live server and a backup server. A live server can have only *one* backup server, and a backup server can be owned by only *one* live server.

When a live server crashes or is brought down in the correct mode, the backup server currently in passive mode will become the new live server. If the new live server is configured to allow automatic failback, it will detect the old live server coming back up and automatically stop, allowing the old live server to start receiving messages again.



### NOTE

If you deploy just one pair of live / backup servers, you cannot effectively use a load balancer in front of the pair because the backup instance is not actively processing messages. Moreover, services such as JNDI and the Undertow web server are not active on the backup server either. For these reasons, deploying JEE applications to an instance of JBoss EAP being used as a backup messaging server is not supported.

### 30.2. HA POLICIES

JBoss EAP messaging supports two different strategies for backing up a server: replication and shared store. Use the **ha-policy** attribute of the **server** configuration element to assign the policy of your choice to the given server. There are four valid values for **ha-policy**:

- **replication-master**
- **replication-slave**
- **shared-store-master**
- **shared-store-slave**

As you can see, the value specifies whether the server uses a [data replication](#) or a [shared store](#) ha policy, and whether it takes the role of master or slave.

Use the management CLI console to add an **ha-policy** to the server of your choice.



## NOTE

The examples below assume you are running JBoss EAP using the **standalone-full-ha** configuration profile.

```
/subsystem=messaging-activemq/server=<SERVER>/ha-policy=<POLICY>:add
```

For example, use the following command to add the **replication-master** policy to the **default** server.

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-master:add
```

The **replication-master** policy is configured with the default values. Values to override the default configuration can be included when you add the policy. The management CLI command to read the current configuration uses the following basic syntax.

```
/subsystem=messaging-activemq/server=<SERVER>/ha-policy=<POLICY>:read-resource
```

For example, use the following command to read the current configuration for the **replication-master** policy that was added above to the **default** server. The output is also included to highlight the default configuration.

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-master:read-resource

{
  "outcome" => "success",
  "result" => {
    "check-for-live-server" => false,
    "cluster-name" => undefined,
    "group-name" => undefined,
    "initial-replication-sync-timeout" => 30000L
  }
}
```

See [Data Replication](#) and [Shared Store](#) for details on the configuration options available for each policy.

## 30.3. DATA REPLICATION


When using replication, the live and the backup server pairs do not share the same data directories, all data synchronization is done over the network. Therefore all (persistent) data received by the live server will be duplicated to the backup.

If the live server is cleanly shut down, the backup server will activate and clients will failover to backup. This behavior is pre-determined and is therefore not configurable when using data replication.

The backup server will first need to synchronize all existing data from the live server before replacing it. Unlike shared storage, therefore, a replicating backup will not be fully operational immediately after startup. The time it will take for the synchronization to happen depends on the amount of data to be synchronized and the network speed. Also note that clients are blocked for the duration of **initial-replication-sync-timeout** when the backup is started. After this timeout elapses, clients will be unblocked, even if synchronization is not completed.

After a successful failover, the backup's journal will start holding newer data than the data on the live server. You can configure the original live server to perform a failback and become the live server once restarted. A failback will synchronize data between the backup and the live server before the live server comes back online.

In cases where both servers are shut down, the administrator will have to determine which server's journal has the latest data. If the backup journal has the latest data, copy that journal to the live server. Otherwise, whenever it activates again, the backup will replicate the stale journal data from the live server and will delete its own journal data. If the live server's data is the latest, no action is needed and the servers can be started normally.



### IMPORTANT

Due to higher latencies and a potentially unreliable network between data centers, the configuration and use of replicated journals for high availability between data centers is neither recommended nor supported.

The replicating live and backup pair must be part of a cluster, and it is the **cluster-connection** configuration element that defines how a backup server will find its live pair. See [Configuring Cluster Connections](#) for details on how to configure a cluster connection. When configuring a **cluster-connection**, keep in mind the following:

- Both the live and backup server must be part of the same cluster. Notice that even a simple live/backup replicating pair will require a cluster configuration.
- Their cluster user and password must match.

Specify a pair of live / backup servers by configuring the **group-name** attribute in either the master or the slave element. A backup server will only connect to a live server that shares the same **group-name**.

As a simple example of using a **group-name**, suppose you have 2 live servers and 2 backup servers. Because each live server needs to pair with its own backup, you would assign names like so:

- live1 and backup1 will both have a **group-name** of **pair1**.
- live2 and backup2 will both have a **group-name** of **pair2**.

In the above example, server **backup1** will search for the live server with the same **group-name**, **pair1**, which in this case is the server **live1**.

Much like in the shared store case, when the live server stops or crashes, its replicating, paired backup will become active and take over its duties. Specifically, the paired backup will become active when it loses connection to its live server. This can be problematic because this can also happen because of a temporary network problem. In order to address this issue, the paired backup will try to determine whether it still can connect to the other servers in the cluster. If it can connect to more than half the servers, it will become active. If it loses communication to its live server plus more than half the other

servers in the cluster, the paired backup will wait and try reconnecting with the live server. This reduces the risk of a "split brain" situation where both the backup and live servers are processing messages without the other knowing it.



### IMPORTANT

This is an important distinction from a shared store backup, where the backup will activate and start to serve client requests if it does not find a live server and the file lock on the journal was released. Note also that in replication the backup server does not know whether any data it might have is up to date, so it really cannot decide to activate automatically. To activate a replicating backup server using the data it has, the administrator must change its configuration to make it a live server by changing slave to master.

### 30.3.1. Configuring Data Replication

Below are two examples showing the basic configuration for both a live and a backup server residing in the cluster named **my-cluster** and in the backup group named **group1**.

The steps below use the management CLI to provide a basic configuration for both a live and a backup server residing in the cluster named **my-cluster** and in the backup group named **group1**.



### NOTE

The examples below assume you are running JBoss EAP using the **standalone-full-ha** configuration profile.

#### Management CLI Commands to Configure a Live Server for Data Replication

1. Add the **ha-policy** to the Live Server

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-
master:add(check-for-live-server=true,cluster-name=my-cluster,group-
name=group1)
```

The **check-for-live-server** attribute tells the live server to check to make sure that no other server has its given id within the cluster.

2. Add the **ha-policy** to the Backup Server

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-
slave:add(cluster-name=my-cluster,group-name=group1)
```

3. Confirm a shared **cluster-connection** exists.

Proper communication between the live and backup servers requires a **cluster-connection**. Use the following management CLI command to confirm that the same **cluster-connection** is configured on both the live and backup servers. The example uses the default **cluster-connection** found in the **standalone-full-ha** configuration profile, which should be sufficient for most use cases. See [Configuring Cluster Connections](#) for details on how to configure a cluster connection.

Use the following management CLI command to confirm that both the live and backup server are using the same cluster-connection.

-

```
/subsystem=messaging-activemq/server=default/cluster-connection=my-cluster:read-resource
```

If the **cluster-connection** exists, the output will provide the current configuration. Otherwise an error message will be displayed.

See [All Replication Configuration](#) for details on all configuration attributes.

### 30.3.2. All Replication Configuration

You can use the management CLI to add configuration to a policy after it has been added. The commands to do so follow the basic syntax below.

```
/subsystem=messaging-activemq/server=default/ha-policy=<POLICY>:write-attribute(name=<ATTRIBUTE>, value=<VALUE>)
```

For example, to set the value of the **restart-backup** attribute to **true**, use the following command.

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-slave:write-attribute(name=restart-backup, value=true)
```

The following tables provide the HA configuration attributes for the **replication-master** node and **replication-slave** configuration elements.

**Table 30.1. Attributes for replication-master**

Attribute	Description
check-for-live-server	Set to <b>true</b> to tell this server to check the cluster for another server using the same server ID when starting up. Default is <b>false</b> .
cluster-name	Name of the cluster used for replication.
group-name	If set, backup servers will only pair with live servers with the matching <b>group-name</b> .
initial-replication-sync-timeout	How long to wait in milliseconds until the initiation replication is synchronized. Default is <b>30000</b> .

**Table 30.2. Attributes for replication-slave**

Attribute	Description
-----------	-------------

Attribute	Description
allow-failback	Whether this server will automatically stop when another places a request to take over its place. A typical use case is when live server requests to resume active processing after a restart or failure recovery. A backup server with <b>allow-failback</b> set to <b>true</b> would yield to the live server once it rejoined the cluster and requested to resume processing. Default is <b>true</b> .
cluster-name.	Name of the cluster used for replication.
group-name.	If set, backup servers will pair only with live servers with the matching <b>group-name</b> .
initial-replication-sync-timeout.	How long to wait in milliseconds until the initiation replication is synchronized. Default is <b>30000</b> .
max-saved-replicated-journal-size.	Specifies how many times a replicated backup server can restart after moving its files on start. After reaching the maximum, the server will stop permanently after if fails back. Default is <b>2</b> .
restart-backup.	Set to <b>true</b> to tell this backup server to restart once it has been stopped because of failback. Default is <b>true</b> .

### 30.3.3. Preventing Cluster Connection Timeouts

Each live and backup pair uses a **cluster-connection** to communicate. The **call-timeout** attribute of a **cluster-connection** sets the amount of a time a server will wait for a response after making a call to another server on the cluster. The default value for **call-timeout** is 30 seconds, which is sufficient for most use cases. However, there are situations where the backup server might be unable to process replication packets coming from the live server. This may happen, for example, when the initial pre-creation of journal files takes too much time, due to slow disk operations or to a large value for **journal-min-files**. If timeouts like this occur you will see a line in your logs similar to the one below.

```
AMQ222207: The backup server is not responding promptly introducing
latency beyond the limit. Replication server being disconnected now.
```



#### WARNING

If a line like the one above appears in your logs that means that the replication process has stopped. You must restart the backup server to reinitiate replication.



To prevent cluster connection timeouts, consider the following options:

- Increase the **call-timeout** of the **cluster-connection**. See [Configuring Cluster Connections](#) for more information.
- Decrease the value of **journal-min-files**. See [Configuring Persistence](#) for more information.

### 30.3.4. Removing Old Journal Directories

A backup server will move its journals to a new location when it starts to synchronize with a live server. By default the journal directories are located in **data/activemq** directory under **EAP\_HOME/standalone**. For domains, each server will have its own **serverX/data/activemq** directory located under **EAP\_HOME/domain/servers**. The directories are named **bindings**, **journal**, **largemessages** and **paging**. See [Configuring Persistence](#) and [Configuring Paging](#) for more information about these directories.

Once moved, the new locations are always given a digit suffix. For example, on the first synchronization the journal directories will be moved to **bindings1**, **journal1**, **largemessages1** and **paging1**. If another synchronization starts due to a new failover then the suffix for the "moved" directories will be increased by 1. The original directories will store the data synchronized from the live server.

By default a backup server is configured to manage two occurrences of failing over and failing back. After the second failback, the backup server will stop and will no longer synchronize with the live server. You will need to remove old messaging directories before the server will manage another failover and restart the backup server.

You can change the number of managed failovers using the **max-saved-replicated-journal-size** attribute on the backup server. Setting **max-saved-replicated-journal-size** to **-1** will keep an unlimited number of backup journals. This setting can cause disk space issues, however, since there is no auto-deletion of the oldest journal directories. Consequently, the oldest directories will remain on disk and will need to be removed.



#### NOTE

Since JBoss EAP 7.0.2, you no longer need to manually remove old directories because there is a new mechanism that automatically does it. However, there was an issue that prevented the automatic clean up process from triggering because the **max-saved-replicated-journal-size** on the live server was set to **-1**. This issue was fixed in Red Hat JBoss Enterprise Application Platform 7.0 Update 05. The fix sets the **max-saved-replicated-journal-size** on the live server to **2**. The property will remain read-only.

You can remove old messaging directories on the backup server either manually or as part of a routine cron job. The safest option is to remove the oldest directories by creation timestamp. For example, after experiencing a second failover, **data/activemq** might contain the following **journal** directories:

```
drwxr-xr-x  May 20 18:12 journal
drwxr-xr-x  May 10 18:12 journal1
drwxr-xr-x  May 15 18:12 journal2
```

In this case you should delete only **journal1**. You **should not** remove **journal** and **journal2**. **journal** contains the files the server has currently, and **journal2** holds your last chance to recover old messages, if, for example, the new journal was created by accident.

Do the same for the other directories, **bindings**, **largemessages** and **paging**, if they are present.

### Old Directories on the Live Server

There are scenarios in which the live server, and not the backup, will be the server that contains old journal directories. When these scenarios occur, it is important to remember that the number of old journals on the live server can grow indefinitely because, unlike the backup server, the live server does not use the **max-saved-replicated-journal-size** configuration option. Therefore, it is very important to prune the old directories before they significantly impact disk capacity. The method to remove old journals from the live server is the same as it is for the backup server. However, there is no need to restart the live server after the old directories are deleted.

One scenario that would cause the live server to increase the number of journal directories is when it comes back online after the backup server had already begun to process messages. Once restarted, the live server will synchronize with the backup server, so that a failback to the live server can happen. As part of the synchronization, the live server will move its outdate **journal** directory to **journal1** so that **journal1** can store up-to-date data from the backup server.

### 30.3.5. Updating Dedicated Live and Backup Servers

If the live and backup servers are deployed in a dedicated topology, where each server is running in its own instance of JBoss EAP, follow the steps below to ensure a smooth update and restart of the cluster.

1. Cleanly shut down the backup servers.
2. Cleanly shut down the live servers.
3. Update the configuration of the live and backup servers.
4. Start the live servers.
5. Start the backup servers.

### 30.3.6. Limitations of Data Replication: Split Brain Processing

A "split brain" situation occurs when both a live server and its backup are active at the same time. Both servers can serve clients and process messages without the other knowing it. In this situation there is no longer any message replication between the live and backup servers. A split situation can happen if there is network failure between the two servers.

For example, if the connection between a live server and a network router is broken, the backup server will lose the connection to the live server. However, because the backup can still connect to more than half the servers in the cluster, it becomes active. Recall that a backup will also activate if there is just one live-backup pair and the backup server loses connectivity to the live server. When both servers are active within the cluster, two undesired situations can happen:

1. Remote clients fail over to the backup server, but local clients such as MDBs will use the live server. Both nodes will have completely different journals, resulting in split brain processing.
2. The broken connection to the live server is fixed after remote clients have already failed over to the backup server. Any new clients will be connected to the live server while old clients continue to use the backup, which also results in a split brain scenario.

Customers should implement a reliable network between each pair of live and backup servers to reduce the risk of split brain processing when using data replication. For example, use duplicated Network Interface Cards and other network redundancies.

## 30.4. SHARED STORE

This style of high availability differs from data replication in that it requires a shared file system which is accessible by both the live and backup node. This means that the server pairs use the same location for their [paging](#), [message journal](#), [bindings journal](#), and [large messages](#) in their configuration.



### NOTE

Using a shared store is not supported on Windows. It is supported on Red Hat Enterprise Linux when using Red Hat versions of GFS2 or NFSv4. In addition, GFS2 is supported only with an ASYNCIO journal type, while NFSv4 is supported with both ASYNCIO and NIO journal types.


Also, each participating server in the pair, live and backup, will need to have a **cluster-connection** defined, even if not part of a cluster, because the **cluster-connection** defines how the backup server announces its presence to its live server and any other nodes. See [Configuring Cluster Connections](#) for details on how this is done.

When failover occurs and a backup server takes over, it will need to load the persistent storage from the shared file system before clients can connect to it. This style of high availability differs from data replication in that it requires a shared file system which is accessible by both the live and backup pair. Typically this will be some kind of high performance Storage Area Network, or SAN. Red Hat does not recommend using Network Attached Storage, known as a NAS, for your storage solution.


The advantage of shared store high availability is that no replication occurs between the live and backup nodes, this means it does not suffer any performance penalties due to the overhead of replication during normal operation.

The disadvantage of shared store replication is that when the backup server activates it needs to load the journal from the shared store which can take some time depending on the amount of data in the store. Also, it requires a shared storage solution supported by JBoss EAP.

If you require the highest performance during normal operation, Red Hat recommends having access to a highly performant SAN and accept the slightly slower failover costs. Exact costs will depend on the amount of data.



### 30.4.1. Configuring a Shared Store

**NOTE**

The examples below assume you are running JBoss EAP using the **standalone-full-ha** configuration profile.

1. Add the **ha-policy** to the Live Server.

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-master:add
```

2. Add the **ha-policy** to the Backup Server.

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-slave:add
```

3. Confirm a shared **cluster-connection** exists.

Proper communication between the live and backup servers requires a **cluster-connection**. Use the following management CLI command to confirm that the same **cluster-connection** is configured on both the live and backup servers. The example uses the default **cluster-connection** found in the **standalone-full-ha** configuration profile, which should be sufficient for most use cases. See [Configuring Cluster Connections](#) for details on how to configure a cluster connection.

```
/subsystem=messaging-activemq/server=default/cluster-connection=my-cluster:read-resource
```

If the **cluster-connection** exists, the output will provide the current configuration. Otherwise an error message will be displayed.

See [All Shared Store Configuration](#) for details on all configuration attributes for shared store policies.

### 30.4.2. All Shared Store Configuration

Use the management CLI to add configuration to a policy after it has been added. The commands to do so follow the basic syntax below.

```
/subsystem=messaging-activemq/server=default/ha-policy=<POLICY>:write-attribute(name=<ATTRIBUTE>, value=<VALUE>)
```

For example, to set the value of the **restart-backup** attribute to **true**, use the following command.

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-slave:write-attribute(name=restart-backup, value=true)
```

**Table 30.3. Attributes of the shared-store-master configuration element.**

Attribute	Description
failover-on-server-shutdown	Set to <b>true</b> to tell this server to failover when it is normally shut down. Default is <b>false</b> .

**Table 30.4. Attributes of the shared-store-slave configuration element.**

Attribute	Description
allow-failback	Set to <b>true</b> to tell this server to automatically stop when another places a request to take over its place. The use case is when a regular server stops and its backup takes over its duties, later the main server restarts and requests the server (the former backup) to stop operating. Default is <b>true</b> .
failover-on-server-shutdown	Set to <b>true</b> to tell this server to failover when it is normally shut down. Default is <b>false</b> .
restart-backup	Set to <b>true</b> to tell this server to restart once it has been stopped because of failback or scaling down. Default is <b>true</b> .

## 30.5. FAILING BACK TO A LIVE SERVER

After a live server has failed and a backup taken has taken over its duties, you may want to restart the live server and have clients fail back to it.

In case of a shared store, simply restart the original live server and kill the new live server by killing the process itself. Alternatively, you can set **allow-fail-back** to **true** on the slave which will force it to automatically stop once the master is back online. The management CLI command to set **allow-fail-back** looks like the following:

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-slave:write-attribute(name=allow-fail-back,value=true)
```

In replication HA mode you need to set an extra property **check-for-live-server** to **true** in the master configuration.

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-master:write-attribute(name=check-for-live-server,value=true)
```

If set to **true**, a live server will search the cluster during startup for another server using its nodeID. If it finds one, it will contact this server and try to "fail-back". Since this is a remote replication scenario, the original live server will have to synchronize its data with the backup running with its ID. Once they are in sync, it will request the backup server to shut down so it can take over active processing. This behavior allows the original live server to determine whether there was a fail-over, and if so whether the server that took its duties is still running or not.

**WARNING**

Be aware that if you restart a live server after failover to backup has occurred then `check-for-live-server` must be set to `true`. If not then the live server will start at once without checking that its backup server is running. This will result in situation in which the live and backup are running at the same time causing the delivery of duplicate messages to all newly connected clients.

For shared stores, it is also possible to cause failover to occur on normal server shut down, to enable this set `failover-on-server-shutdown` to `true` in the HA configuration on either the master or slave like so:

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-slave:write-attribute(name=failover-on-server-shutdown,value=true)
```

You can also force the running backup server to shut down when the original live server comes back up, allowing the original live server to take over automatically, by setting `allow-failback` to `true`.

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-slave:write-attribute(name=allow-failback,value=true)
```

## 30.6. COLOCATED BACKUP SERVERS

JBoss EAP also makes it possible to colocate backup messaging servers in the same JVM as another live server. Take for example a simple two node cluster of standalone servers where each live server colocates the backup for the other. You can use either a shared store or a replicated HA policy when colocating servers in this way. There are two important things to remember when configuring messaging servers for colocation.

First, each `server` element in the configuration will need its own `remote-connector` and `remote-acceptor` that listen on unique ports. For example, a live server can be configured to listen on port `5445`, while its backup uses `5446`. The ports are defined in `socket-binding` elements that must be added to the default `socket-binding-group`. Cluster-related configuration elements in each `server` configuration will use the new `remote-connector`. The relevant configuration is included in each of the examples that follow.

**NOTE**

`http-acceptors` and `http-connectors` can be used to create colocated HA topology too, but only if different `http-listeners` are used for the master and slave connectors. For this reason, using `remote-acceptors` and `remote-connectors` is recommended.

Secondly, remember to properly configure paths for journal related directories. For example, in a shared store colocated topology, both the live server and its backup, colocated on another live server, must be configured to share directory locations for the `binding` and `message` journals, for `large messages`, and for `paging`.

### 30.6.1. Configuring Colocated Servers

The example management CLI commands used in the steps below illustrate how to configure a simple two node cluster employing a colocated topology. Examples are included for both a shared store and a data replication HA policy.



## NOTE

The examples below assume you are running JBoss EAP using the **full-ha** configuration profile.

Steps to configure a two node colocated cluster. A live and backup server will live on each node. The colocated backup on node1 is paired with the live server colocated on node 2, and the backup server on node 2 will be paired with the live server on node 1.

1. Add socket bindings for all servers.

The colocated messaging servers will need separate ports to use for communication. Use the management CLI to create a **socket-binding** for each colocated server. In the example below, port **5445** is used by the live server, and port **5446** will be used by the backup. Note that the new bindings are added to the default **socket-binding-group** named **standard-sockets**.

```
/socket-binding-group=standard-sockets/socket-binding=messaging:add(port=5445)
```

```
/socket-binding-group=standard-sockets/socket-binding=messaging-backup:add(port=5446)
```

2. Modify the default server on each instance to use an HA policy

The default server on each node will become the live server. Use the management CLI to add the preferred HA policy.

If using a shared store, the command is:

```
/subsystem=messaging-activemq/server=default/ha-policy=shared-store-master:add
```

For data replication, the default server on each node one should be configured with a unique **group-name**. In this example, the first command below is executed on node one, and the second on node two.

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-master:add(cluster-name=my-cluster,group-name=group1,check-for-live-server=true)
```

```
/subsystem=messaging-activemq/server=default/ha-policy=replication-master:add(cluster-name=my-cluster,group-name=group2,check-for-live-server=true)
```

3. Colocate a new backup server with each live server.

Add a new server to each instance of JBoss EAP that will colocate with the default, live, server. The new server will backup the default server on the other node. Use the following management CLI command to create a new server named **backup**.

```
/subsystem=messaging-activemq/server=backup:add
```

Next, configure the new server to use the HA policy of choice.

If using a shared store, the command is:

```
/subsystem=messaging-activemq/server=backup/ha-policy=shared-store-slave:add
```

For data replication, the backup servers should be configured with the **group-name** of the live server on the other node. In this example, the first command below is executed on node one, and the second on node two.

```
/subsystem=messaging-activemq/server=backup/ha-policy=replication-slave:add(cluster-name=my-cluster,group-name=group2)
```

```
/subsystem=messaging-activemq/server=backup/ha-policy=replication-slave:add(cluster-name=my-cluster,group-name=group1)
```

#### 4. Configure the directory locations for all servers.

Now that the servers are configured for HA, each must have the appropriately configured locations for the [binding](#) and [message](#) journals, for [large messages](#), and for [paging](#) if it is being used.

Use the commands below to configure the directory locations for all servers.

For a shared store, the **path** values for the live server on node one should point to the same location on a supported file system as the backup server on node two. Likewise for the live server on node two and its backup on node one.

Configure the directory locations for node one:

```
/subsystem=messaging-activemq/server=default/path=bindings-directory:write-attribute(name=path,value=/path/to/shared/bindings-A)
```

```
/subsystem=messaging-activemq/server=default/path=journal-directory:write-attribute(name=path,value=/path/to/shared/journal-A)
```

```
/subsystem=messaging-activemq/server=default/path=large-messages-directory:write-attribute(name=path,value=/path/to/shared/largemessages-A)
```

```
/subsystem=messaging-activemq/server=default/path=paging-directory:write-attribute(name=path,value=/path/to/shared/paging-A)
```

```
/subsystem=messaging-activemq/server=backup/path=bindings-directory:write-attribute(name=path,value=/path/to/shared/bindings-B)
```

```
/subsystem=messaging-activemq/server=backup/path=journal-directory:write-attribute(name=path,value=/path/to/shared/journal-B)
```

```
/subsystem=messaging-activemq/server=backup/path=large-messages-directory:write-attribute(name=path,value=/path/to/shared/largemessages-B)
```



```
/subsystem=messaging-activemq/server=backup/path=paging-
directory:write-attribute(name=path,value=/path/to/shared/paging-B)
```

Configure the directory locations for node two:

```
/subsystem=messaging-activemq/server=default/path=bindings-
directory:write-attribute(name=path,value=/path/to/shared/bindings-
B)

/subsystem=messaging-activemq/server=default/path=journal-
directory:write-attribute(name=path,value=/path/to/shared/journal-B)

/subsystem=messaging-activemq/server=default/path=large-messages-
directory:write-
attribute(name=path,value=/path/to/shared/largemessages-B)

/subsystem=messaging-activemq/server=default/path=paging-
directory:write-attribute(name=path,value=/path/to/shared/paging-B)

/subsystem=messaging-activemq/server=backup/path=bindings-
directory:write-attribute(name=path,value=/path/to/shared/bindings-
A)

/subsystem=messaging-activemq/server=backup/path=journal-
directory:write-attribute(name=path,value=/path/to/shared/journal-A)

/subsystem=messaging-activemq/server=backup/path=large-messages-
directory:write-
attribute(name=path,value=/path/to/shared/largemessages-A)

/subsystem=messaging-activemq/server=backup/path=paging-
directory:write-attribute(name=path,value=/path/to/shared/paging-A)
```

For data replication, each server will use its own directories and will not share them with any other server. In the example commands below, each value for a **path** location is assumed to be a unique location on a file system. In fact, there is no need to change the directory locations for the live servers since they will use the default locations. The backup servers will still need to be configured with unique locations however.

Configure the directory locations for node one:

```
/subsystem=messaging-activemq/server=backup/path=bindings-
directory:write-attribute(name=path,value=activemq/bindings-B)

/subsystem=messaging-activemq/server=backup/path=journal-
directory:write-attribute(name=path,value=activemq/journal-B)

/subsystem=messaging-activemq/server=backup/path=large-messages-
directory:write-attribute(name=path,value=activemq/largemessages-B)

/subsystem=messaging-activemq/server=backup/path=paging-
directory:write-attribute(name=path,value=activemq/paging-B)
```

Configure the directory locations for node two:

```

/subsystem=messaging-activemq/server=backup/path=bindings-
directory:write-attribute(name=path,value=activemq/bindings-B)

/subsystem=messaging-activemq/server=backup/path=journal-
directory:write-attribute(name=path,value=activemq/journal-B)

/subsystem=messaging-activemq/server=backup/path=large-messages-
directory:write-attribute(name=path,value=activemq/largemessages-B)

/subsystem=messaging-activemq/server=backup/path=paging-
directory:write-attribute(name=path,value=activemq/paging-B)

```

5. Add a new acceptor and connector to all servers.

Each messaging server will need to be configured with a new remote-connector and a new remote-acceptor that will use the socket-bindings created in the first step. This step allows a server to receive and send communication over the port established in the binding. The examples below add a remote-acceptor and remote-connector to the default server. To do the same for the backup server, replace **server=default** with **server=backup**.

Add the acceptor:

```

/subsystem=messaging-activemq/server=default/remote-
acceptor=netty:add(socket-binding=messaging)

/subsystem=messaging-activemq/server=backup/remote-
acceptor=netty:add(socket-binding=messaging-backup)

```

Add the connector:

```

/subsystem=messaging-activemq/server=default/remote-
connector=netty:add(socket-binding=messaging)

/subsystem=messaging-activemq/server=backup/remote-
connector=netty:add(socket-binding=messaging-backup)

```

6. Configure the **cluster-connection** for all servers.

Each messaging server will need a **cluster-connection** and a broadcast-group for proper communication. Use the following management CLI commands to configure these elements. Note that the **cluster-connection** must be added to the backup servers, with the configuration passed in.

```

/subsystem=messaging-activemq/server=default/cluster-connection=my-
cluster:write-attribute(name=connector-name,value=netty)

/subsystem=messaging-activemq/server=backup/cluster-connection=my-
cluster:add(connector-name=netty,cluster-connection-address=jms)

```

7. Configure the **broadcast-group** for all servers.

The default server already has a **broadcast-group**. Configure it so that it uses only the new connector. First, clear any connectors that may be already pre-configured as part of the **full-ha** profile. Then, add the **netty** connector.

```

/subsystem=messaging-activemq/server=default/broadcast-group=bg-

```

```
group1:list-clear(name=connectors)
```

```
/subsystem=messaging-activemq/server=default/broadcast-group=bg-  
group1:list-add(name=connectors,value="netty")
```

Backup servers just need to have the **netty** connector added since they are not pre-configured.

```
/subsystem=messaging-activemq/server=backup/broadcast-group=bg-  
group1:add(connectors=[netty],jgroups-channel=activemq-cluster)
```

8. Modify the default **RemoteConnectionFactory** on the live servers.

The default servers already have a **connection-factory** used for remote connections. Configure it so that it uses only the new **remote-connector**. First, clear any connectors that may be already pre-configured as part of the **full-ha** profile. Then, add the **netty** connector.

```
/subsystem=messaging-activemq/server=default/connection-  
factory=RemoteConnectionFactory:list-clear(name=connectors)
```

```
/subsystem=messaging-activemq/server=default/connection-  
factory=RemoteConnectionFactory:list-  
add(name=connectors,value=netty)
```

9. Add a **discovery-group** to the backup servers.

Backup servers use a **discovery-group** for initial cluster enrollment. Add one to each backup server using the command below. Note the reference to the **socket-binding** created in Step 1.

```
/subsystem=messaging-activemq/server=backup/discovery-group=dg-  
group-backup:add(jgroups-channel=activemq-cluster)
```

## 30.7. FAILOVER MODES

JBoss EAP messaging defines two types of client failover:

- Automatic client failover
- Application-level client failover

JBoss EAP messaging also provides 100% transparent automatic reattachment of connections to the same server (e.g. in case of transient network problems). This is similar to failover, except it is reconnecting to the same server and is discussed in [Client Reconnection and Session Reattachment](#).

During failover, if the client has consumers on any non persistent or temporary queues, those queues will be automatically recreated during failover on the backup node, since the backup node will not have any knowledge of non persistent queues.

### 30.7.1. Automatic Client Failover

JBoss EAP messaging clients can be configured to receive knowledge of all live and backup servers, so that in the event of a connection failure at the client - live server connection, the client will detect the failure and reconnect to the backup server. The backup server will then automatically recreate any sessions and consumers that existed on each connection before failover, thus saving the user from having to hand-code manual reconnection logic.

A JBoss EAP messaging client detects connection failure when it has not received packets from the server within the time given by **client-failure-check-period** as explained in [Detecting Dead Connections](#).

If the client does not receive data in the allotted time, it will assume the connection has failed and attempt failover. If the socket is closed by the operating system, the server process might be killed rather than the server hardware itself crashing for example, the client will failover straight away.

JBoss EAP messaging clients can be configured to discover the list of live-backup server pairs in a number of different ways. They can be configured with explicit endpoints, for example, but the most common way is for the client to receive information about the cluster topology when it first connects to the cluster. See [Server Discovery](#) for more information.

The default HA configuration includes a **cluster-connection** that uses the recommended **http-connector** for cluster communication. This is the same **http-connector** that remote clients use when making connections to the server using the default **RemoteConnectionFactory**. While it is not recommended, you can use a different connector. If you use your own connector, make sure it is included as part of the configuration for both the **connection-factory** to be used by the remote client and the **cluster-connection** used by the cluster nodes. See [Configuring the Messaging Transports](#) and [Cluster Connections](#) for more information on connectors and cluster connections.



### WARNING

The **connector** defined in the **connection-factory** to be used by a JMS client must be the same one defined in the **cluster-connection** used by the cluster. Otherwise, the client will not be able to update its topology of the underlying live/backup pairs and therefore will not know the location of the backup server.

Use CLI commands to review the configuration for both the **connection-factory** and the **cluster-connection**. For example, to read the current configuration for the **connection-factory** named **RemoteConnectionFactory** use the following command.

```
/subsystem=messaging-activemq/server=default/connection-  
factory=RemoteConnectionFactory:read-resource
```

Likewise, the command below reads the configuration for the **cluster-connection** named **my-cluster**.

```
/subsystem=messaging-activemq/server=default/cluster-connection=my-  
cluster:read-resource
```

To enable automatic client failover, the client must be configured to allow non-zero reconnection attempts. See [Client Reconnection and Session Reattachment](#) for more information. By default, failover will occur only after at least one connection has been made to the live server. In other words, failover will not occur if the client fails to make an initial connection to the live server. If it does fail its initial attempt, a client would simply retry connecting to the live server according to the **reconnect-attempts** property and fail after the configured number of attempts.

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=reconnect-
attempts,value=<NEW_VALUE>)
```

An exception to this rule is the case where there is only one pair of live - backup servers, and no other live server, and a remote MDB is connected to the live server when it is cleanly shut down. If the MDB has configured `@ActivationConfigProperty(propertyName = "rebalanceConnections", propertyValue = "true")`, it tries to rebalance its connection to another live server and will not failover to the backup.

### Failing Over on the Initial Connection

Since the client does not learn about the full topology until after the first connection is made, there is a window of time where it does not know about the backup. If a failure happens at this point the client can only try reconnecting to the original live server. To configure how many attempts the client will make you can set the property `initialConnectAttempts` on the `ClientSessionFactoryImpl` or `ActiveMQConnectionFactory`.

Alternatively in the server configuration, you can set the `initial-connect-attempts` attribute of the connection factory used by the client. The default for this is `0`, that is, try only once. Once the number of attempts has been made, an exception will be thrown.

```
/subsystem=messaging-activemq/server=default/connection-
factory=RemoteConnectionFactory:write-attribute(name=initial-connect-
attempts,value=<NEW_VALUE>)
```

### About Server Replication

JBoss EAP messaging does not replicate full server state between live and backup servers. When the new session is automatically recreated on the backup, it won't have any knowledge of the messages already sent or acknowledged during that session. Any in-flight sends or acknowledgements at the time of failover may also be lost.

By replicating full server state, JBoss EAP messaging could theoretically provide a 100% transparent seamless failover, avoiding any lost messages or acknowledgements. However, doing so comes at a great cost: replicating the full server state, including the queues and session. This would require replication of the entire server state machine. That is, every operation on the live server would have to be replicated on the replica servers in the exact same global order to ensure a consistent replica state. This is extremely hard to do in a performant and scalable way, especially considering that multiple threads are changing the live server state concurrently.

It is possible to provide full state machine replication using techniques such as virtual synchrony, but this does not scale well and effectively serializes all operations to a single thread, dramatically reducing concurrency. Other techniques for multi-threaded active replication exist such as replicating lock states or replicating thread scheduling, but this is very hard to achieve at a Java level.

Consequently, it was not worth reducing performance and concurrency for the sake of 100% transparent failover. Even without 100% transparent failover, it is simple to guarantee once and only once delivery, even in the case of failure, by using a combination of duplicate detection and retrying of transactions. However this is not 100% transparent to the client code.

#### 30.7.1.1. Handling Blocking Calls During Failover

If the client code is in a blocking call to the server, i.e. it is waiting for a response to continue its execution, during a failover, the new session will not have any knowledge of the call that was in progress. The blocked call might otherwise hang forever, waiting for a response that will never come.

To prevent this, JBoss EAP messaging will unblock any blocking calls that were in progress at the time of failover by making them throw a `javax.jms.JMSEException`, if using JMS, or an `ActiveMQException` with error code `ActiveMQException.UNBLOCKED` if using the core API. It is up to the client code to catch this exception and retry any operations if desired.

If the method being unblocked is a call to `commit()`, or `prepare()`, then the transaction will be automatically rolled back and JBoss EAP messaging will throw a `javax.jms.TransactionRolledBackException`, if using JMS, or a `ActiveMQException` with error code `ActiveMQException.TRANSACTION_ROLLED_BACK` if using the core API.

### 30.7.1.2. Handling Failover With Transactions

If the session is transactional and messages have already been sent or acknowledged in the current transaction, then the server cannot be sure whether messages or acknowledgements were lost during the failover.

Consequently the transaction will be marked as rollback-only, and any subsequent attempt to commit it will throw a `javax.jms.TransactionRolledBackException`, if using JMS, or a `ActiveMQException` with error code `ActiveMQException.TRANSACTION_ROLLED_BACK` if using the core API.



#### WARNING

The caveat to this rule is when XA is used either via JMS or through the core API. If a two phase commit is used and `prepare()` has already been called then rolling back could cause a `HeuristicMixedException`. Because of this the commit will throw a `XAException.XA_RETRY` exception. This informs the Transaction Manager that it should retry the commit at some later point in time, a side effect of this is that any non persistent messages will be lost. To avoid this from happening, be sure to use persistent messages when using XA. With acknowledgements this is not an issue since they are flushed to the server before `prepare()` gets called.

It is up to the user to catch the exception and perform any client side local rollback code as necessary. There is no need to manually rollback the session since it is already rolled back. The user can then just retry the transactional operations again on the same session.

If failover occurs when a commit call is being executed, the server, as previously described, will unblock the call to prevent a hang, since no response will come back. In this case it is not easy for the client to determine whether the transaction commit was actually processed on the live server before failure occurred.



#### NOTE

If XA is being used either via JMS or through the core API then an `XAException.XA_RETRY` is thrown. This is to inform Transaction Managers that a retry should occur at some point. At some later point in time the Transaction Manager will retry the commit. If the original commit has not occurred, it will still exist and be committed. If it does not exist, then it is assumed to have been committed, although the transaction manager may log a warning.

To remedy this, the client can enable duplicate detection in the transaction, and retry the transaction operations again after the call is unblocked. See [Duplicate Message Detection](#) for information on how detection is configured on the server. If the transaction had indeed been committed on the live server successfully before failover, duplicate detection will ensure that any durable messages resent in the transaction will be ignored on the server to prevent them getting sent more than once when the transaction is retried.

### 30.7.1.3. Getting Notified of Connection Failure

JMS provides a standard mechanism for sending asynchronously notifications of a connection failure: `java.jms.ExceptionListener`. Please consult the [JMS javadoc](#) for more information on this class. The core API also provides a similar feature in the form of the class `org.apache.activemq.artemis.core.client.SessionFailureListener`.

Any `ExceptionListener` or `SessionFailureListener` instance will always be called by JBoss EAP in case of a connection failure, whether the connection was successfully failed over, reconnected, or reattached. However, you can find out if the reconnect or reattach has happened by inspecting the value for the `failedOver` flag passed into `connectionFailed()` on `SessionFailureListener` or the error code on the `javax.jms.JMSEException` which will be one of the following:

JMSEException error codes

Error code	Description
FAILOVER	Failover has occurred and we have successfully reattached or reconnected.
DISCONNECT	No failover has occurred and we are disconnected.

### 30.7.2. Application-Level Failover

In some cases you may not want automatic client failover, and prefer to handle any connection failure yourself, and code your own manually reconnection logic in your own failure handler. We define this as application-level failover, since the failover is handled at the user application level.

To implement application-level failover if you're using JMS set an `ExceptionListener` class on the JMS connection. The `ExceptionListener` will be called by JBoss EAP messaging in the event that connection failure is detected. In your `ExceptionListener`, you would close your old JMS connections, potentially look up new connection factory instances from JNDI and creating new connections.

If you are using the core API, then the procedure is very similar: you would set a `FailureListener` on the core `ClientSession` instances.

## 30.8. DETECTING DEAD CONNECTIONS

This section discusses connection time to live (TTL) and explains how JBoss EAP messaging handles crashed clients and clients that have exited without cleanly closing their resources.

### Cleaning up Dead Connection Resources on the Server

Before a JBoss EAP client application exits, it should close its resources in a controlled manner, using a `finally` block.

Below is an example of a core client appropriately closing its session and session factory in a **finally** block:

```
ServerLocator locator = null;
ClientSessionFactory sf = null;
ClientSession session = null;

try {
    locator = ActiveMQClient.createServerLocatorWithoutHA(..);

    sf = locator.createClientSessionFactory();

    session = sf.createSession(...);

    ... do some stuff with the session...
}
finally {
    if (session != null) {
        session.close();
    }

    if (sf != null) {
        sf.close();
    }

    if (locator != null) {
        locator.close();
    }
}
```

And here is an example of a well behaved JMS client application:

```
Connection jmsConnection = null;

try {
    ConnectionFactory jmsConnectionFactory =
ActiveMQJMSClient.createConnectionFactoryWithoutHA(...);

    jmsConnection = jmsConnectionFactory.createConnection();

    ... do some stuff with the connection...
}
finally {
    if (connection != null) {
        connection.close();
    }
}
```

Unfortunately sometimes clients crash and do not have a chance to clean up their resources. If this occurs, it can leave server side resources hanging on the server. If these resources are not removed they would cause a resource leak on the server, and over time this likely would result in the server running out of memory or other resources.

When looking to clean up dead client resources, it is important to be aware of the fact that sometimes the network between the client and the server can fail and then come back, allowing the client to reconnect.



Because JBoss EAP supports client reconnection, it is important that it not clean up "dead" server side resources too soon, or clients will be prevented any client from reconnecting and regaining their old sessions on the server.

JBoss EAP makes all of this configurable. For each **ClientSessionFactory** configured, a Time-To-Live, or TTL, property can be used to set how long the server will keep a connection alive in milliseconds in the absence of any data from the client. The client will automatically send "ping" packets periodically to prevent the server from closing its connection. If the server does not receive any packets on a connection for the length of the TTL time, it will automatically close all the sessions on the server that relate to that connection.

If you are using JMS, the connection TTL is defined by the **ConnectionTTL** attribute on a **ActiveMQConnectionFactory** instance, or if you are deploying JMS connection factory instances direct into JNDI on the server side, you can specify it in the xml config, using the parameter **connectionTtl**.

The default value for **ConnectionTTL** on an network-based connection, such as an **http-connector**, is **60000**, i.e. 1 minute. The default value for connection TTL on a internal connection, e.g. an **in-vm** connection, is **-1**. A value of **-1** for **ConnectionTTL** means the server will never time out the connection on the server side.

If you do not want clients to specify their own connection TTL, you can set a global value on the server side. This can be done by specifying the **connection-ttl-override** attribute in the server configuration. The default value for **connection-ttl-override** is **-1** which means "do not override", i.e. let clients use their own values.

### Closing Core Sessions or JMS Connections

It is important that all core client sessions and JMS connections are always closed explicitly in a **finally** block when you are finished using them.

If you fail to do so, JBoss EAP will detect this at garbage collection time. It will then close the connection and log a warning similar to the following:

```
[Finalizer] 20:14:43,244 WARNING
[org.apache.activemq.artemis.core.client.impl.DelegatingSession] I'm
closing a ClientSession you left open. Please make sure you close all
ClientSessions explicitly before let
ting them go out of scope!
[Finalizer] 20:14:43,244 WARNING
[org.apache.activemq.artemis.core.client.impl.DelegatingSession] The
session you didn't close was created here:
java.lang.Exception
    at org.apache.activemq.artemis.core.client.impl.DelegatingSession.
<init>(DelegatingSession.java:83)
    at org.acme.yourproject.YourClass (YourClass.java:666)
```

Note that if you are using JMS the warning will involve a JMS connection, not a client session. Also, the log will tell you the exact line of code where the unclosed JMS connection or core client session was instantiated. This will enable you to pinpoint the error in your code and correct it appropriately.

### Detecting Failure from the Client Side

As long as the client is receiving data from the server it will consider the connection to be alive. If the client does not receive any packets for **client-failure-check-period** milliseconds, it will consider the connection failed and will either initiate failover, or call any **FailureListener** instances, or

**ExceptionListener** instances if you are using JMS, depending on how the client has been configured.

If you are using JMS the behavior is defined by the **ClientFailureCheckPeriod** attribute on a **ActiveMQConnectionFactory** instance.

The default value for client failure check period on a network connection, for example an HTTP connection, is **30000**, or 30 seconds. The default value for client failure check period on an in-vm connection, is **-1**. A value of **-1** means the client will never fail the connection on the client side if no data is received from the server. Whatever the type of connection, the check period is typically much lower than the value for connection TTL on the server so that clients can reconnect in case of transitory failure.

### Configuring Asynchronous Connection Execution

Most packets received on the server side are executed on the **remoting** thread. These packets represent short-running operations and are always executed on the **remoting** thread for performance reasons.

However, by default some kinds of packets are executed using a thread from a thread pool so that the **remoting** thread is not tied up for too long. Please note that processing operations asynchronously on another thread adds a little more latency. These packets are:

```
org.apache.activemq.artemis.core.protocol.core.impl.wireformat.RollbackMessage
org.apache.activemq.artemis.core.protocol.core.impl.wireformat.SessionCloseMessage
org.apache.activemq.artemis.core.protocol.core.impl.wireformat.SessionCommitMessage
org.apache.activemq.artemis.core.protocol.core.impl.wireformat.SessionXCommitMessage
org.apache.activemq.artemis.core.protocol.core.impl.wireformat.SessionXPrepareMessage
org.apache.activemq.artemis.core.protocol.core.impl.wireformat.SessionXRollbackMessage
```

To disable asynchronous connection execution, set the parameter **async-connection-execution-enabled** to **false**. The default value is **true**.

## 30.9. CLIENT RECONNECTION AND SESSION REATTACHMENT

JBoss EAP messaging clients can be configured to automatically reconnect or reattach to the server in the event that a failure is detected in the connection between the client and the server.

### Transparent Session Reattachment

If the failure was due to some transient cause such as a temporary network outage, and the target server was not restarted, the sessions will still exist on the server, assuming the client has not been disconnected for more than the value of **connection-ttl**. See [Detecting Dead Connections](#).

In this scenario, JBoss EAP will automatically reattach the client sessions to the server sessions when the re-connection is made. This is done 100% transparently and the client can continue exactly as if nothing had happened.

As JBoss EAP messaging clients send commands to their servers they store each sent command in an in-memory buffer. When a connection fails and the client subsequently attempts to reattach to the same server, as part of the reattachment protocol, the server gives the client the id of the last command it successfully received.

If the client has sent more commands than were received before failover it can replay any sent commands from its buffer so that the client and server can reconcile their states.

The size in bytes of this buffer is set by the **confirmationWindowSize** property. When the server has received **confirmationWindowSize** bytes of commands and processed them it will send back a command confirmation to the client, and the client can then free up space in the buffer.

If you are using the JMS service on the server to load your JMS connection factory instances into JNDI, then this property can be configured in the server configuration, by setting the **confirmation-window-size** attribute of the chosen **connection-factory**. If you are using JMS but not using JNDI then you can set these values directly on the **ActiveMQConnectionFactory** instance using the appropriate setter method, **setConfirmationWindowSize**. If you are using the core API, the **ServerLocator** instance has a **setConfirmationWindowSize** method exposed as well.

Setting **confirmationWindowSize** to **-1**, which is also the default, disables any buffering and prevents any reattachment from occurring, forcing a reconnect instead.

### Session Reconnection

Alternatively, the server might have actually been restarted after crashing or it might have been stopped. In such a case any sessions will no longer exist on the server and it will not be possible to 100% transparently reattach to them.

In this case, JBoss EAP will automatically reconnect the connection and recreate any sessions and consumers on the server corresponding to the sessions and consumers on the client. This process is exactly the same as what happens when failing over to a backup server.

Client reconnection is also used internally by components such as core bridges to allow them to reconnect to their target servers.

See the section on [Automatic Client Failover](#) to get a full understanding of how transacted and non-transacted sessions are reconnected during a reconnect and what you need to do to maintain once and only once delivery guarantees.

### Configuring Reconnection Attributes

Client reconnection is configured by setting the following properties:

- **retryInterval**. This optional parameter sets the period in milliseconds between subsequent reconnection attempts, if the connection to the target server has failed. The default value is **2000** milliseconds.
- **retryIntervalMultiplier**. This optional parameter sets a multiplier to apply to the time since the last retry to compute the time to the next retry. This allows you to implement an exponential backoff between retry attempts.  
For example, if you set **retryInterval** to **1000** ms and set **retryIntervalMultiplier** to **2.0**, then, if the first reconnect attempt fails, the client will wait **1000** ms then **2000** ms then **4000** ms between subsequent reconnection attempts.

The default value is **1.0** meaning each reconnect attempt is spaced at equal intervals.

- **maxRetryInterval**. This optional parameter sets the maximum retry interval that will be used. When setting **retryIntervalMultiplier** it would otherwise be possible that subsequent retries exponentially increase to ridiculously large values. By setting this parameter you can set an upper limit on that value. The default value is **2000** milliseconds.
- **reconnectAttempts**. This optional parameter sets the total number of reconnect attempts to make before giving up and shutting down. A value of **-1** signifies an unlimited number of attempts. The default value is **0**.

If you are using JMS and JNDI on the client to look up your JMS connection factory instances then you can specify these parameters in the JNDI context environment. For example, your **jndi.properties** file might look like the following.

```
java.naming.factory.initial = ActiveMQInitialContextFactory
connection.ConnectionFactory=tcp://localhost:8080?
retryInterval=1000&retryIntervalMultiplier=1.5&maxRetryInterval=60000&reco
nnectAttempts=1000
```

If you are using JMS, but instantiating your JMS connection factory directly, you can specify the parameters using the appropriate setter methods on the **ActiveMQConnectionFactory** immediately after creating it.

If you are using the core API and instantiating the **ServerLocator** instance directly you can also specify the parameters using the appropriate setter methods on the **ServerLocator** immediately after creating it.

If your client does manage to reconnect but the session is no longer available on the server, for instance if the server has been restarted or it has timed out, then the client will not be able to reattach, and any **ExceptionListener** or **FailureListener** instances registered on the connection or session will be called.

### **ExceptionListeners and SessionFailureListeners**

Note that when a client reconnects or reattaches, any registered JMS **ExceptionListener** or core API **SessionFailureListener** will be called.

## CHAPTER 31. RESOURCE ADAPTERS

A Java Connector Architecture (JCA) Resource Adapter lets your applications communicate with any messaging provider. It configures how JEE components such as MDBs and other EJBs, and even Servlets, can send or receive messages.

### 31.1. ABOUT THE INTEGRATED ARTEMIS RESOURCE ADAPTER

JBoss EAP 7 includes an integrated Artemis resource adapter, which uses the **pooled-connection-factory** element to configure the outbound and inbound connections of the resource adapter.

#### Outbound Connection

Outbound connections are defined using the **pooled-connection-factory** element, which is then used in Java EE deployments by EJBs and servlets to send messages to and receive messages from queues or topics. Because connections created from connection factories are created in the scope of the application server, they can use application server features like the following:

- Connection pooling
- Authentication using the security domains defined by the application server
- Participation in XA transactions using the transaction manager

This is a major difference with a **pooled-connection-factory** as these features are not available with a basic **connection-factory** like **InVmConnectionFactory** or **RemoteConnectionFactory**. Also, be aware that with a connection factory defined using **pooled-connection-factory**, it is not possible to do a lookup using JNDI from an external standalone JMS client.

#### Inbound Connections

Inbound connections are used only by message-driven beans (MDBs) to receive message from a queue or a topic. MDBs are stateless session beans that listen on a queue or topic. They must implement the public **onMessage(Message message)** method, which is called when a message is sent to a queue or a topic. The Artemis resource adapter is responsible for receiving the message from the queue or the topic and passing it to the **onMessage(Message message)** method. For this purpose it configures the inbound connection, which defines the location of the integrated Artemis server and some additional elements.

Each MDB session bean uses a thread from the client thread pool to consume the message from the destination. The default maximum thread pool size is set to eight (8) times the number of CPU cores. For systems with many MDB sessions, such as test suites, this can potentially lead to thread exhaustion and force MDBs to wait for a free thread from the pool. You can configure the maximum client thread pool size using the **activemq.artemis.client.global.thread.pool.max.size** system property. To set the maximum thread pool size to **128**, pass the following argument on command line when starting the server.

```
-Dactivemq.artemis.client.global.thread.pool.max.size=128
```

For more information about MDBs, see [Message Driven Beans](#) in *Developing EJB Applications* for JBoss EAP.

### 31.2. USING THE INTEGRATED ARTEMIS RESOURCE ADAPTER FOR REMOTE CONNECTIONS

JBoss EAP includes a resource adapter to make connections to its integrated ActiveMQ Artemis messaging server. By default the **pooled-connection-factory** defined in the **messaging-activemq** subsystem uses the adapter to make the connections. However, you can use the same resource adapter to make connections to an Artemis server running inside a remote instance of JBoss EAP as well.

To connect to an Artemis server running inside a remote instance of JBoss EAP, create a new **pooled-connection-factory** by following the steps below.

1. Create an outbound-socket-binding pointing to the remote messaging server:

```
/socket-binding-group=standard-sockets/remote-destination-outbound-socket-binding=remote-server:add(host=<server host>, port=8080)
```

2. Create a remote-connector referencing the outbound-socket-binding created in step 1.

```
/subsystem=messaging-activemq/server=default/http-connector=remote-http-connector:add(socket-binding=remote-server, endpoint=http-acceptor)
```

3. Create a pooled-connection-factory referencing the remote-connector created in step 2.

```
/subsystem=messaging-activemq/server=default/pooled-connection-factory=remote-artemis:add(connectors=[remote-http-connector], entries=[java:/jms/remoteCF])
```

### Configuring an MDB to use a pooled-connection-factory

After the **pooled-connection-factory** is configured to connect to a remote Artemis server, Message-Driven Beans (MDB) wanting to read messages from the remote server must be annotated with the **@ResourceAdapter** annotation using the name of the **pooled-connection-factory** resource.

```
import org.jboss.ejb3.annotation.ResourceAdapter;

@ResourceAdapter("remote-artemis")
@MessageDriven(name = "MyMDB", activationConfig = { ... })
public class MyMDB implements MessageListener {
    public void onMessage(Message message) {
        ...
    }
}
```

If the MDB needs to send messages to the remote server, it must inject the **pooled-connection-factory** by looking it up using one of its **JNDI entries**.

```
@Inject
@JMSConnectionFactory("java:/jms/remoteCF")
private JMSContext context;
```

### Configuring the JMS destination

An MDB must also specify the destination from which it will consume messages. The standard way to do this is to define a **destinationLookup** activation config property that corresponds to a JNDI lookup on the local server.

```

@ResourceAdapter("remote-artemis")
@MessageDriven(name = "MyMDB", activationConfig = {
    @ActivationConfigProperty(propertyName = "destinationLookup",
        propertyValue = "myQueue"),
    ...
})
public class MyMDB implements MessageListener {
    ...
}

```

If the local server does not include a JNDI binding for the remote Artemis server, specify the name of the destination, as configured in the remote Artemis server, using the **destination** activation config property and set the **useJNDI** activation config property to **false**. This instructs the Artemis resource adapter to automatically create the JMS destination without requiring a JNDI lookup.

```

@ResourceAdapter("remote-artemis")
@MessageDriven(name = "MyMDB", activationConfig = {
    @ActivationConfigProperty(propertyName = "useJNDI", propertyValue =
        "false"),
    @ActivationConfigProperty(propertyName = "destination", propertyValue
        = "myQueue"),
    ...
})
public class MyMDB implements MessageListener {
    ...
}

```

In the above example, the activation config properties configure the MDB to consume messages from the JMS Queue named **myQueue** hosted on the remote Artemis server. In most cases, the MDB does not need to lookup other destinations to process the consumed messages, and it can use the **JMSReplyTo** destination if it is defined on the message.

If the MDB needs any other JMS destinations defined on the remote server, it must use client-side JNDI. See [Connecting to a Server](#) for more information.

## 31.3. DEPLOYING AN A-MQ RESOURCE ADAPTER

You can deploy the resource adapter provided by the Red Hat JBoss A-MQ product and have, for example, Red Hat JBoss A-MQ 6.2.0, become the external JMS provider for JBoss EAP.

See the [Red Hat JBoss A-MQ documentation](#) for details on how to deploy and configure an A-MQ resource adapter.

### 31.3.1. Issues with the A-MQ 6 Resource Adapter

- JBoss EAP will track and monitor applications, looking for unclosed resources. While useful in many cases, such monitoring might cause unexpected behavior when an application tries to re-use a closed instance of **UserTransaction** in a single method. Add the attribute **tracking="false"** to the **<connection-definition/>** element when configuring the A-MQ resource adapter if your applications re-use connections in this way.

```

<connection-definition class-name="..." tracking="false" ... />

```

- The A-MQ 6 resource adapter does not implement **XAResourceWrapper** from the Narayana



API, which is used by JBoss EAP. Consequently, when the Transaction Manager sends a commit to all the XA transaction participants and then crashes while waiting for a reply, it will go on indefinitely logging warnings until records of the committed transaction are removed from its object store.

- The A-MQ 6 resource adapter returns the code **XAER\_RMERR** when an error, such as a network disconnection, occurs during the call of the commit method protocol. This behavior breaks the XA specification since the correct return code should be **XAER\_RMFAIL** or **XAER\_RETRY**. Consequently, the transaction is left in an unknown state on the message broker side, which can cause data inconsistency in some cases. A message will be logged similar to the one below when the unexpected error code is returned.

```
WARN [com.arjuna.ats.jtax] ...: XAResourceRecord.rollback caused an
XA error: ARJUNA016099: Unknown error code:0 from resource ... in
transaction ...: javax.transaction.xa.XAException: Transaction ...
has not been started.
```

- Red Hat JBoss A-MQ 6.x supports the JMS 1.1 specification that is included with Java EE 6. It does not support the JMS 2.0 specification that is included with Java EE 7 and supported in JBoss EAP 7. If you need to send messages to a remote A-MQ broker, you must use the [JMS 1.1 API](#) within your application code. For more information about JBoss A-MQ supported standards, see [Red Hat JBoss A-MQ Supported Standards and Protocols](#).

## 31.4. DEPLOYING THE IBM WEBSHERE® MQ RESOURCE ADAPTER

### About IBM WebSphere® MQ

IBM WebSphere® MQ is the Messaging Oriented Middleware (MOM) product offering from IBM that allows applications on distributed systems to communicate with each other. This is accomplished through the use of messages and message queues. IBM WebSphere® MQ is responsible for delivering messages to the message queues and for transferring data to other queue managers using message channels. For more information about IBM WebSphere® MQ, see [IBM MQ](#).

### Summary

This section covers the steps to deploy and configure the IBM WebSphere® MQ resource adapter in JBoss EAP. JBoss EAP 7 is certified with IBM WebSphere® MQ 7.5.0.4. This deployment and configuration can be accomplished by manually editing configuration files, using the management CLI tool, or using the web-based management console.

### Prerequisites

Before you get started, you must verify the version of the IBM WebSphere® MQ resource adapter and understand some of its configuration properties.

- The IBM WebSphere® MQ resource adapter is supplied as a Resource Archive (RAR) file called **wmq.jmsra.rar**. You can obtain the **wmq.jmsra.rar** file from **/opt/mqm/java/lib/jca/wmq.jmsra.rar**. You must use version 7.5.0.x. See the note above for information about the required version.
- You must know the values of the following IBM WebSphere® MQ configuration properties. Refer to the [IBM WebSphere® MQ documentation](#) for your version of the product for details about these properties.
  - MQ.QUEUE.MANAGER: The name of the IBM WebSphere® MQ queue manager
  - MQ.HOST.NAME: The host name used to connect to the IBM WebSphere® MQ queue manager



- MQ.CHANNEL.NAME: The server channel used to connect to the IBM WebSphere® MQ queue manager
- MQ.QUEUE.NAME: The name of the destination queue
- MQ.TOPIC.NAME: The name of the destination topic
- MQ.PORT: The port used to connect to the IBM WebSphere® MQ queue manager
- MQ.CLIENT: The transport type
- For outbound connections, you must also be familiar with the following configuration property:
  - MQ.CONNECTIONFACTORY.NAME: The name of the connection factory instance that will provide the connection to the remote system

## Procedure to Deploy the IBM WebSphere® Resource Adapter



### NOTE

The following are default configurations provided by IBM and are subject to change. For more information, refer to [IBM WebSphere® MQ documentation](#) for your version of the product.

1. First, deploy the resource adapter manually by copying the `wmq.jmsra.rar` file to the `EAP_HOME/standalone/deployments/` directory.
2. Next, use the management CLI to add the resource adapter and configure it.

```
/subsystem=resource-adapters/resource-adapter=wmq.jmsra.rar:add(archive=wmq.jmsra.rar, transaction-support=XATransaction)
```

Note that the **transaction-support** element was set to **XATransaction**. When using transactions, be sure to supply the security domain of the XA recovery datasource, as in the example below.

```
/subsystem=resource-adapters/resource-adapter=test/connection-definitions=test:write-attribute(name=recovery-security-domain,value=myDomain)
```

For more information about XA Recovery see [Configuring XA Recovery](#) in the JBoss EAP *Configuration Guide*.

For non-transactional deployments, change the value of **transaction-support** to `NoTransaction`.

```
/subsystem=resource-adapters/resource-adapter=wmq.jmsra.rar:add(archive=wmq.jmsra.rar, transaction-support=NoTransaction)
```

3. Now that the resource adapter is created, you can add the necessary configuration elements to it.
  - a. Add an **admin-object** for queues and configure its properties.

```

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/admin-objects=queue-ao:add(class-
name=com.ibm.mq.connector.outbound.MQQueueProxy, jndi-
name=java:jboss/MQ.QUEUE.NAME)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/admin-objects=queue-ao/config-
properties=baseQueueName:add(value=MQ.QUEUE.NAME)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/admin-objects=queue-ao/config-
properties=baseQueueManagerName:add(value=MQ.QUEUE.MANAGER)

```

- b. Add an **admin-object** for topics and configure its properties.

```

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/admin-objects=topic-ao:add(class-
name=com.ibm.mq.connector.outbound.MQTopicProxy, jndi-
name=java:jboss/MQ.TOPIC.NAME)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/admin-objects=topic-ao/config-
properties=baseTopicName:add(value=MQ.TOPIC.NAME)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/admin-objects=topic-ao/config-
properties=brokerPubQueueManager:add(value=MQ.QUEUE.MANAGER)

```

- c. Add a connection definition for a managed connection factory and configure its properties

```

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/connection-definitions=mq-cd:add(class-
name=com.ibm.mq.connector.outbound.ManagedConnectionFactoryImpl,
jndi-name=java:jboss/MQ.CONNECTIONFACTORY.NAME, tracking=false)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/connection-definitions=mq-cd/config-
properties=hostName:add(value=MQ.HOST.NAME)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/connection-definitions=mq-cd/config-
properties=port:add(value=MQ.PORT)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/connection-definitions=mq-cd/config-
properties=channel:add(value=MQ.CHANNEL.NAME)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/connection-definitions=mq-cd/config-
properties=transportType:add(value=MQ.CLIENT)

/subsystem=resource-adapters/resource-
adapter=wmq.jmsra.rar/connection-definitions=mq-cd/config-
properties=queueManager:add(value=MQ.QUEUE.MANAGER)

```

- If you want to change the default provider for the EJB3 messaging system in JBoss EAP from JBoss EAP 7 messaging to IBM WebSphere® MQ, use the management CLI to modify the **ejb3** subsystem as follows:

```
/subsystem=ejb3:write-attribute(name=default-resource-adapter-name,value=wmq.jmsra.rar)
```

- Configure the **@ActivationConfigProperty** and **@ResourceAdapter** annotations in the MDB code as follows:

```
@MessageDriven(name="WebSphereMQMDB", activationConfig = {
    @ActivationConfigProperty(propertyName =
"destinationType",propertyValue = "javax.jms.Queue"),
    @ActivationConfigProperty(propertyName = "useJNDI",
propertyValue = "false"),
    @ActivationConfigProperty(propertyName = "hostName",
propertyValue = "MQ.HOST.NAME"),
    @ActivationConfigProperty(propertyName = "port", propertyValue =
"MQ.PORT"),
    @ActivationConfigProperty(propertyName = "channel",
propertyValue = "MQ.CHANNEL.NAME"),
    @ActivationConfigProperty(propertyName = "queueManager",
propertyValue = "MQ.QUEUE.MANAGER"),
    @ActivationConfigProperty(propertyName = "destination",
propertyValue = "MQ.QUEUE.NAME"),
    @ActivationConfigProperty(propertyName = "transportType",
propertyValue = "MQ.CLIENT")
})

@ResourceAdapter(value = "wmq.jmsra-VERSION.rar")
@TransactionAttribute(TransactionAttributeType.NOT_SUPPORTED)
public class WebSphereMQMDB implements MessageListener {
}
```

Be sure to replace the "VERSION" with the actual version in the name of the RAR.

### 31.4.1. The Limitations and Known Problems for IBM WebSphere® MQ 7.5 Resource Adapter

- The deployment of the IBM WebSphere® MQ 7.5 resource adapter does not load the **javax.jms.api** module for your deployment. It also does not provide support for the new Java EE 7 annotations like **@JMSConnectionFactoryDefinitions**, **@JMSDestinationDefinition**. It is necessary to have the **messaging-activemq** subsystem in the configuration to enable it. If you do not want the JBoss EAP messaging server to be started, add an empty **messaging-activemq** subsystem.

```
<subsystem xmlns="urn:jboss:domain:messaging-activemq:1.0">
</subsystem>
```

- Calling the **createConnection()** method on the IBM WebSphere® MQ 7.5.0.4 resource adapter implementation of **QueueConnectionFactory** and **TopicConnectionFactory**, with or without parameters, leads to a **javax.jms.JMSEException** to be thrown.

- The IBM WebSphere® MQ 7.5.0.4 resource adapter is able to read messages from queues and topics even before the connection has started.
- If the transaction manager calls second rollback against the same XID, then IBM WebSphere® MQ 7.5 resource adapter does not return **XAException** with error code **XAER\_NOTA** but it rather returns **XAER\_RMFAIL** which is incorrect. This does not have an impact on data consistency but the logged warnings in the server log are impacted.
- The IBM WebSphere® MQ resource adapter does not implement **XAResourceWrapper** from the Narayana API, which is used by JBoss EAP. Consequently, when the Transaction Manager sends a commit to all the XA transaction participants and then crashes while waiting for a reply, it will go on indefinitely logging warnings until records of the committed transaction are removed from its object store.
- There is a known issue in IBM WebSphere® MQ Resource Adapter version 7.5.0.3 and earlier that causes periodic recovery to fail with an XA exception with messages similar to the following in the JBoss EAP server log:

```
WARN [com.arjuna.ats.jta] (Periodic Recovery) ARJUNA016027: Local
XARecoveryModule.xaRecovery got XA exception XAException.XAER_INVALID:
javax.transaction.xa.XAException: The method 'xa_recover' has failed
with errorCode '-5'.
```

A fix is available in later versions, but remember the behavior of version 7.5.0.4 as noted in this list. A detailed description of this issue can be found here: <http://www-01.ibm.com/support/docview.wss?uid=swg11C97579>.

- With IBM WebSphere® MQ 7.5 and later, the **commit()** and **rollback()** methods on **UserTransaction** close any JMS connections that are part of that transaction. Resolve this issue by setting **tracking="false"** in the **connection-definition** as in the following example.

```
/subsystem=resource-adapters/resource-adapter=eis.rar/connection-
definitions=myConnectionDef:write-
attribute(name=tracking,value=false)
```

- The **wmq.jmsra.rar** resource adapter contains jars with dependencies configured to be found in specific locations. When these hard-coded locations are not discovered, JBoss EAP will log warnings like those below. These warnings are harmless from a functional point of view and require no action.

```
14:38:27,543 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry connector.jar in
/content/wmq.jmsra.rar/com.ibm.mq.jar does not point to a valid jar
for a Class-Path reference.
14:38:27,551 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry jta.jar in
/content/wmq.jmsra.rar/com.ibm.mq.jmqi.jar does not point to a
valid jar for a Class-Path reference.
14:38:27,552 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry ldap.jar in
/content/wmq.jmsra.rar/com.ibm.mqjms.jar does not point to a valid
jar for a Class-Path reference.
14:38:27,553 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry jndi.jar in
```

```

/content/wmq.jmsra.rar/com.ibm.mqjms.jar does not point to a valid
jar for a Class-Path reference.
14:38:27,553 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry fscontext.jar in
/content/wmq.jmsra.rar/com.ibm.mqjms.jar does not point to a valid
jar for a Class-Path reference.
14:38:27,553 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry providerutil.jar in
/content/wmq.jmsra.rar/com.ibm.mqjms.jar does not point to a valid
jar for a Class-Path reference.
14:38:27,555 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry jms.jar in
/content/wmq.jmsra.rar/com.ibm.msg.client.jms.jar does not point to
a valid jar for a Class-Path reference.
14:38:27,557 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry rmm.jar in
/content/wmq.jmsra.rar/com.ibm.msg.client.wmq.v6.jar does not point
to a valid jar for a Class-Path reference.
14:38:27,557 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry CL3Export.jar in
/content/wmq.jmsra.rar/com.ibm.msg.client.wmq.v6.jar does not point
to a valid jar for a Class-Path reference.
14:38:27,557 WARN [org.jboss.as.server.deployment] (MSC service
thread 1-8) WFLYSRV0059: Class Path entry CL3Nonexport.jar in
/content/wmq.jmsra.rar/com.ibm.msg.client.wmq.v6.jar does not point
to a valid jar for a Class-Path reference.
...
14:38:28,018 WARN [org.jboss.as.connector.deployers.RADeployer]
(MSC service thread 1-5) IJ020017: Invalid archive:
file:/home/mnovak/tmp/jboss-eap-
7.0/standalone/tmp/vfs/temp/tempc3f7f9d35e24ba16/content-
12e0796e05502d20/contents/
14:38:28,124 WARN [org.jboss.as.connector.deployers.RaXmlDeployer]
(MSC service thread 1-8) IJ020017: Invalid archive:
file:/home/mnovak/tmp/jboss-eap-
7.0/standalone/tmp/vfs/temp/tempc3f7f9d35e24ba16/content-
12e0796e05502d20/contents/

```

## 31.5. DEPLOYING A GENERIC JMS RESOURCE ADAPTER




### WARNING

Generic JMS resource adapters are not supported in JBoss EAP 7.

The Tibco EMS 6 JMS provider that is used in this section to demonstrate how to configure and deploy a generic JMS resource adapter is *not* supported in JBoss EAP 7.

JBoss EAP can be configured to work with third-party JMS providers, however not all JMS providers produce a JMS JCA resource adapter for integration with Java application platforms. This procedure

covers the steps required to configure the generic JMS resource adapter included in JBoss EAP to connect to a JMS provider. In this procedure, Tibco EMS 6.3 is used as an example JMS provider. Other JMS providers may require different configuration.



### IMPORTANT

Before using the generic JMS resource adapter, check with the JMS provider to see if they have their own resource adapter that can be used with JBoss EAP. The generic JMS JCA resource adapter should only be used when a JMS provider does not provide its own resource adapter.

Before you can configure a generic resource adapter, you will need to do the following:

- Your JMS provider server must already be configured and ready for use. Any binaries required for the provider's JMS implementation will be needed.
- You will need to know the values of the following JMS provider properties to be able to look up its JMS resources, such as connection factories, queues or topics.
  - `java.naming.factory.initial`
  - `java.naming.provider.url`
  - `java.naming.factory.url.pkgs`

In the example XML used in this procedure, these parameters are written as **PROVIDER\_FACTORY\_INITIAL**, **PROVIDER\_URL**, and **PROVIDER\_CONNECTION\_FACTORY** respectively. Replace these placeholders with the JMS provider values for your environment.

### 31.5.1. Configure a Generic JMS Resource Adapter for Use with a Third-party JMS Provider

1. Create and configure the resource adapter module.

Create a JBoss EAP module that contains all the libraries required to connect and communicate with the JMS provider. This module will be named `com.tibco.tibjms`.

- Create the following directory structure: **EAP\_HOME/modules/com/tibco/tibjms/main**
- Copy the binaries required for the provider's JMS implementation to **EAP\_HOME/modules/com/tibco/tibjms/main**.



### NOTE

For Tibco EMS, the binaries required are `tibjms.jar` and `tibcrypt.jar` from the Tibco installation's `lib` directory.

- Create a `module.xml` file in **EAP\_HOME/modules/com/tibco/tibjms/main** as below, listing the JAR files from the previous steps as resources:

```
<module xmlns="urn:jboss:module:1.1" name="com.tibco.tibjms">
  <resources>
    <!-- all jars required by the JMS provider, in this case
    Tibco -->
    <resource-root path="tibjms.jar"/>
  </resources>
</module>
```

```

    <resource-root path="tibcrypt.jar"/>
  </resources>
  <dependencies>
    <module name="javax.api"/>
    <module name="javax.jms.api"/>
  </dependencies>
</module>

```

- Add the module to the **ee** subsystem using the following CLI command: `/subsystem=ee:list-add(name=global-modules, value={"name" => "com.tibco.tibjms", "slot" => "main"})`
2. Create and configure a JNDI external context to the JMS provider.  
The JMS resources, such as connection factories and destinations, are looked up in the JMS provider. Add an external context in the JBoss EAP instance so that any *local* lookup for this resource will automatically look up the resource on the remote JMS provider.



## NOTE

In this procedure, `EAP_HOME/standalone/configuration/standalone-full.xml` is used as the JBoss EAP configuration file.

Use the management CLI to create an external JNDI context and include its configuration properties. The properties in the example below should be replaced by the correct value to connect to the remote JMS provider. For example, some JMS providers, such as Tibco EMS, do not support the JNDI **lookup(Name)** method. In these cases, add the **org.jboss.as.naming.lookup.by.string** property with a value of **true** to work around this issue. Check the adapter's documentation for information on required properties and their values.

```

/subsystem=naming/binding="java:global/remoteJMS":add(binding-
type=external-
context,module=my.generic.JMSAdapter,class=javax.naming.InitialConte
xt,environment=
[java.naming.factory.initial=my.initial.Factory,java.naming.provider
.url=localhost:9000,java.naming.factory.url.pkgs=my.initial,org.jbos
s.as.naming.lookup.by.string=true])

```

With the external context configured properly, any JNDI lookup to a resource starting with **java:global/remoteJMS/** will be done on the remote JMS provider. As an example, if a message-driven bean performs a JNDI lookup for **java:global/remoteJMS/Queue1**, the external context will connect to the remote JMS provider and perform a lookup for the **Queue1** resource.

Alternatively, you can make a JNDI lookup to the remote server without using an **external-context** when looking up the JNDI name. To do so, use the CLI to create a new binding that references the **external-context**, as in the example below.

```

/subsystem=naming/binding=java:\:/jms\queue\myQueue:add(binding-
type=lookup, lookup=java:global/remoteJMS/jms/queue/myQueue)

```

In the example above, an application that does a JNDI lookup for **java:/jms/queue/myQueue** will locate the queue named **myQueue** on the remote server.

3. Create the generic JMS resource adapter.



Use the management CLI to create the resource adapter

```
/subsystem=resource-adapters/resource-adapter=com.tibco.tibjms:add(module=com.tibco.tibjms,transaction-support=XATransaction)
```

4. Configure the generic JMS resource adapter.

Use the management CLI to configure the resource adapter's **connection-definition** and other elements.

```
/subsystem=resource-adapters/resource-adapter=com.tibco.tibjms/connection-definitions=tibco-cd:add(class-name=org.jboss.resource.adapter.jms.JmsManagedConnectionFactory,jndi-name=java:/jms/XAQCF)
```

```
/subsystem=resource-adapters/resource-adapter=com.tibco.tibjms/connection-definitions=tibco-cd/config-properties=ConnectionFactory:add(value=XAQCF)
```

```
/subsystem=resource-adapters/resource-adapter=com.tibco.tibjms/connection-definitions=tibco-cd/config-properties=JndiParameters:add(value=java.naming.factory.initial=com.tibco.tibjms.naming.TibjmsInitialContextFactory;java.naming.provider.url=TIBCO_EMS_SERVER_HOST_NAME:PORT)
```

```
/subsystem=resource-adapters/resource-adapter=com.tibco.tibjms/connection-definitions=tibco-cd:write-attribute(name=security-application,value=true)
```

5. Configure the default message-driven bean pool with the generic resource adapter. In **EAP\_HOME/standalone/configuration/standalone-full.xml**, in **<subsystem xmlns="urn:jboss:domain:ejb3:1.5">**, update the **<mdb>** configuration with:

```
<mdb>
  <resource-adapter-ref resource-adapter-name="com.tibco.tibjms"/>
  <bean-instance-pool-ref pool-name="mdb-strict-max-pool"/>
</mdb>
```

6. Finally, use the CLI to remove the default server from the **messaging-activemq** subsystem since it is no longer the messaging provider.

```
/subsystem=messaging-activemq/server=default:remove()
```

Removing the default server will avoid any ClassCastExceptions or similar errors when looking up the destination.

The generic JMS resource adapter is now configured and ready for use. Below is an example of using the resource adapter when creating a new message-driven bean.

### Example Using the Generic Resource Adapter

```
@MessageDriven(name = "HelloWorldQueueMDB", activationConfig = {
    // The generic JMS resource adapter requires the JNDI bindings
    // for the actual remote connection factory and destination
```



```

    @ActivationConfigProperty(propertyName = "connectionFactory",
propertyValue = "java:global/remoteJMS/XAQCF"),
    @ActivationConfigProperty(propertyName = "destination", propertyValue =
"java:global/remoteJMS/Queue1"),
    @ActivationConfigProperty(propertyName = "destinationType",
propertyValue = "javax.jms.Queue"),
    @ActivationConfigProperty(propertyName = "acknowledgeMode",
propertyValue = "Auto-acknowledge" ) })
    public class HelloWorldQueueMDB implements MessageListener {
        public void onMessage(Message message) {
            // called every time a message is received from the _Queue1_ queue on
            the JMS provider.
        }
    }
}

```

### IMPORTANT

When using the generic JMS resource adapter, ensure you set the session to be transacted, to avoid a potential **NullPointerException** error. The error occurs because the generic JMS resource adapter attempts processing of parameters, when the Java EE specification states that they are *not* to be processed. This is accomplished by doing the following: **connection.createSession(true, Session.SESSION\_TRANSACTED);**

You can also use the pooled connection factory from the resource adapter:

```

@Resource(lookup = "java:/jms/XAQCF")
private ConnectionFactory cf;

```

It is not possible to inject a resource from an external context directly but it is possible to inject an external context and then perform a lookup. For example, a lookup for a queue deployed in a Tibco EMS broker would be as follows.

```

@Resource(lookup = "java:global/remoteJMS")
private Context context;
...
Queue queue = (Queue) context.lookup("Queue1")

```

## CHAPTER 32. BACKWARD AND FORWARD COMPATIBILITY

JBoss EAP supports both backward and forward compatibility with legacy versions of JBoss EAP that were using HornetQ as their messaging brokers, such as JBoss EAP 6. These two compatibility modes are provided by the JBoss EAP built-in messaging server, ActiveMQ Artemis, that supports the HornetQ's core protocol.

- **Forward compatibility:** Legacy JMS clients using HornetQ can connect to a JBoss EAP 7 server running ActiveMQ Artemis.
- **Backward compatibility:** JBoss EAP 7 JMS clients using JBoss EAP messaging can connect to the legacy JBoss EAP 6 server running HornetQ.

### 32.1. FORWARD COMPATIBILITY

Forward compatibility requires no code changes to legacy JBoss EAP 6 JMS clients. Support is provided by the JBoss EAP **messaging-activemq** subsystem and its resources. To enable support of forward compatibility make the following changes to the configuration of the JBoss EAP 7 server. Example management CLI commands for a standalone server are provided for each step.

- Create a **socket-binding** that listens on port 4447 for remote legacy clients.

```
/socket-binding-group=standard-sockets/socket-binding=legacy-remoting:add(port=4447)
```

- Create a legacy **remote-connector** that will use the **socket-binding** created in the previous step. This is required for JNDI lookups.

```
/subsystem=remoting/connector=legacy-remoting-connector:add(socket-binding=legacy-remoting)
```

- Set up a legacy messaging **socket-binding** that listens on port 5445.

```
/socket-binding-group=standard-sockets/socket-binding=legacy-messaging:add(port=5445)
```

- Set up a **remote-connector** and a **remote-acceptor** in the **messaging-activemq** subsystem that use the binding from the previous step.

```
/subsystem=messaging-activemq/server=default/remote-connector=legacy-messaging-connector:add(socket-binding=legacy-messaging)
```

```
/subsystem=messaging-activemq/server=default/remote-acceptor=legacy-messaging-acceptor:add(socket-binding=legacy-messaging)
```

- Create a legacy HornetQ JMS ConnectionFactory in the **legacy-connection-factory** element of the **messaging-activemq** subsystem.

```
/subsystem=messaging-activemq/server=default/legacy-connection-factory=legacy-discovery:add(entries=[java:jboss/exported/jms/LegacyRemoteConnectionFactory], connectors=[legacy-messaging-connector])
```

- Create legacy HornetQ JMS destinations and include **legacy-entries** attributes to the **jms-queue** or **jms-topic** resources.

```
jms-queue add --queue-address=myQueue --entries=
[java:jboss/exported/jms/myQueue-new] --legacy-entries=
[java:jboss/exported/jms/myQueue]

jms-topic add --topic-address=myTopic --entries=
[java:jboss/exported/jms/myTopic-new] --legacy-entries=
[java:jboss/exported/jms/myTopic]
```

You can add **legacy-entries** to an existing queue or topic by following the below example.

```
/subsystem=messaging-activemq/server=default/jms-
queue=myQueue:write-attribute(name=legacy-entries,value=
[java:jboss/exported/jms/myQueue])
```

While the **entries** attributes are used by JBoss EAP messaging JMS clients, the **legacy-entries** are used by the legacy HornetQ JMS clients. Legacy JMS clients look up this legacy JMS resource to communicate with JBoss EAP 7.



#### NOTE

To avoid any code change in the legacy JMS clients, the legacy JNDI entries configured in the **messaging-activemq** subsystem must match the lookup expected by the legacy JMS client.

### Management CLI migrate Operation

When you run the management CLI **migrate** operation to update your **messaging** subsystem configuration, if the boolean argument **add-legacy-entries** is set to **true**, the **messaging-activemq** subsystem creates the **legacy-connection-factory** resource and adds **legacy-entries** to the **jms-queue** and **jms-topic** resources. The legacy entries in the migrated **messaging-activemq** subsystem will correspond to the entries specified in the legacy **messaging** subsystem and the regular entries are created with a **-new** suffix.

If the boolean argument **add-legacy-entries** is set to **false** when you run the **migrate** operation, no legacy resources are created in the **messaging-activemq** subsystem and legacy JMS clients will not be able to communicate with the JBoss EAP 7 servers.

## 32.2. BACKWARD COMPATIBILITY

Backward compatibility requires no configuration change in the legacy JBoss EAP 7 servers. JBoss EAP 7 JMS clients do not look up resources on the legacy server, but instead use client-side JNDI to create JMS resources. JBoss EAP 7 JMS clients can then use these resources to communicate with the legacy server using the HornetQ core protocol.

**WARNING**

JBoss EAP 7 client connections to a JBoss EAP 5 server are currently not supported.

JBoss EAP messaging supports client-side JNDI to create JMS **ConnectionFactory** and **Destination** resources.

For example, if a JBoss EAP 7 JMS client wants to communicate with a legacy server using a JMS queue named "myQueue", it must use the following properties to configure its JNDI **InitialContext**:

```
java.naming.factory.initial=org.apache.activemq.artemis.jndi.ActiveMQInitialContextFactory
connectionFactory.jms/ConnectionFactory=tcp://<legacy server address>:5445? \

protocolManagerFactoryStr=org.apache.activemq.artemis.core.protocol.hornetq.client.HornetQClientProtocolManagerFactory
queue.jms/myQueue=myQueue
```

The client can then use the **jms/ConnectionFactory** name to create the JMS **ConnectionFactory** and use the **jms/myQueue** to create the JMS **Queue**. Note that the property **protocolManagerFactoryStr=org.apache.activemq.artemis.core.protocol.hornetq.client.HornetQClientProtocolManagerFactory** is mandatory when specifying the URL of the legacy connection factory. This allows the JBoss EAP messaging JMS client to communicate with the HornetQ broker in the legacy server.

## PART IV. PERFORMANCE TUNING

## CHAPTER 33. TUNING JMS

If you use the JMS API, review the following information for tips on how to improve performance.

- **Disable the message ID.**  
If you do not need message IDs, disable them by using the `setDisableMessageID()` method on the `MessageProducer` class. Setting the value to `true` eliminates the overhead of creating a unique ID and decreases the size of the message.
- **Disable the message timestamp.**  
If you do not need message timestamps, disable them by using the `setDisableMessageTimeStamp()` method on the `MessageProducer` class. Setting the value to `true` eliminates the overhead of creating the timestamp and decreases the size of the message.
- **Avoid using `ObjectMessage`.**  
`ObjectMessage` is used to send a message that contains a serialized object, meaning the body of the message, or payload, is sent over the wire as a stream of bytes. The Java serialized form of even small objects is quite large and takes up a lot of space on the wire. It is also slow when compared to custom marshalling techniques. Use `ObjectMessage` only if you cannot use one of the other message types, for example, if you do not know the type of the payload until runtime.
- **Avoid `AUTO_ACKNOWLEDGE`.**  
The choice of acknowledgement mode in a consumer impacts performance because of the additional overhead and traffic incurred by sending the acknowledgment message sent over the network. `AUTO_ACKNOWLEDGE` incurs this overhead because it requires an acknowledgement to be sent from the server for each message received on the client. If you can, use `DUPS_OK_ACKNOWLEDGE`, which acknowledges messages in a lazy manner, `CLIENT_ACKNOWLEDGE`, meaning the client code will call a method to acknowledge the message, or batch up many acknowledgements with one acknowledge or commit in a transacted session.
- **Avoid durable messages.**  
By default, JMS messages are durable. If you do not need durable messages, set them to be **non-durable**. Durable messages incur a lot of overhead because they are persisted to storage.
- **Use `TRANSACTIONED_SESSION` mode to send and receive messages in a single transaction.**  
By batching messages in a single transaction, the ActiveMQ Artemis server integrated in JBoss EAP requires only one network round trip on the commit, not on every send or receive.

---

## CHAPTER 34. TUNING PERSISTENCE

- Put the message journal on its own physical volume.  
One of the advantages of an append-only journal is that disk head movement is minimized. This advantage is lost if the disk is shared. When multiple processes, such as a transaction coordinator, databases, and other journals, read and write from the same disk, performance is impacted because the disk head must skip around between different files. If you are using paging or large messages, make sure they are also put on separate volumes.
- Tune the **journal-min-files** value.  
Set the **journal-min-files** parameter to the number of files that fits your average sustainable rate. If you frequently see new files being created on the journal data directory, meaning a lot data is being persisted, you need to increase the minimal number of files. This allows the journal to reuse, rather than create, new data files.
- Optimize the journal file size.  
The journal file size must be aligned to the capacity of a cylinder on the disk. The default value of **10MB** should be enough on most systems.
- Use the **AIO** journal type.  
For Linux operating systems, keep your journal type as **AIO**. **AIO** scales better than Java **NIO**.
- Tune the **journal-buffer-timeout** value.  
Increasing the **journal-buffer-timeout** value results in increased throughput at the expense of latency.
- Tune the **journal-max-io** value.  
If you are using **AIO**, you might be able improve performance by increasing the **journal-max-io** parameter value. Do not change this value if you are using **NIO**.

## CHAPTER 35. OTHER TUNING OPTIONS

This section describes other places in JBoss EAP messaging that can be tuned.

- Use asynchronous send acknowledgements.  
If you need to send non-transactional, durable messages and do not need a guarantee that they have reached the server by the time the call to `send()` returns, do not set them to be sent blocking. Instead use asynchronous send acknowledgements to get your send acknowledgements returned in a separate stream. However, in the case of a server crash, some messages might be lost.
- Use **pre-acknowledge** mode.  
With **pre-acknowledge** mode, messages are acknowledged before they are sent to the client. This reduces the amount of acknowledgment traffic on the wire. However, if that client crashes, messages will not be redelivered if the client reconnects.
- Disable security.  
There is a small performance boost when you disable security by setting the **security-enabled** attribute to **false**.
- Disable persistence.  
You can turn off message persistence altogether by setting **persistence-enabled** to **false**.
- Sync transactions lazily.  
Setting **journal-sync-transactional** to **false** provides better transactional persistent performance at the expense of some possibility of loss of transactions on failure.
- Sync non-transactional lazily.  
Setting **journal-sync-non-transactional** to **false** provides better non-transactional persistent performance at the expense of some possibility of loss of durable messages on failure.
- Send messages non-blocking.  
To avoid waiting for a network round trip for every message sent, set **block-on-durable-send** and **block-on-non-durable-send** to **false** if you are using JMS and JNDI, or set it directly on the `ServerLocator` by calling the `setBlockOnDurableSend()` and `setBlockOnNonDurableSend()` methods.
- Optimize the **consumer-window-size**.  
If you have very fast consumers, you can increase the **consumer-window-size** to effectively disable consumer flow control.
- Use the core API instead of the JMS API.  
JMS operations must be translated into core operations before the server can handle them, resulting in lower performance than when you use the core API. When using the core API, try to use methods that take **SimpleString** as much as possible. **SimpleString**, unlike `java.lang.String`, does not require copying before it is written to the wire, so if you reuse **SimpleString** instances between calls, you can avoid some unnecessary copying. Note that the core API is not portable to other brokers.



## CHAPTER 36. AVOIDING ANTI-PATTERNS

- Reuse connections, sessions, consumers, and producers where possible.  
The most common messaging anti-pattern is the creation of a new connection, session, and producer for every message sent or consumed. These objects take time to create and may involve several network round trips, so it is a poor use of resources. Always reuse them.



### NOTE

Some popular libraries such as the Spring JMS Template use these anti-patterns. If you are using the Spring JMS Template, you may see poor performance. The Spring JMS Template can only safely be used in an application server which caches JMS sessions, for example, using JCA, and only then for sending messages. It cannot safely be used for synchronously consuming messages, even in an application server.

- Avoid fat messages.  
Verbose formats such as XML take up a lot of space on the wire and performance suffers as result. Avoid XML in message bodies if you can.
- Do not create temporary queues for each request.  
This common anti-pattern involves the temporary queue request-response pattern. With the temporary queue request-response pattern, a message is sent to a target, and a reply-to header is set with the address of a local temporary queue. When the recipient receives the message, they process it, and then send back a response to the address specified in the reply-to header. A common mistake made with this pattern is to create a new temporary queue on each message sent, which drastically reduces performance. Instead, the temporary queue should be reused for many requests.
- Do not use message driven beans unless it is necessary.  
Using MDBs to consume messages is slower than consuming messages using a simple JMS message consumer.

## APPENDIX A. REFERENCE MATERIAL

### A.1. ADDRESS SETTING ATTRIBUTES

Table A.1. Address Setting Attributes

Name	Description
address-full-policy	Determines what happens when an address where <b>max-size-bytes</b> is specified becomes full. Accepted values are <b>PAGE</b> , <b>DROP</b> , <b>FAIL</b> or <b>BLOCK</b> . If the value is <b>PAGE</b> then further messages will be paged to disk. If the value is <b>DROP</b> then further messages will be silently dropped. If the value is <b>FAIL</b> then the messages will be dropped and the client message producers will receive an exception. If the value is <b>BLOCK</b> then client message producers will block when they try and send further messages. <b>PAGE</b> is the default. See <a href="#">About Paging</a> for details on paging.
auto-create-jms-queues	Determines whether the JBoss EAP should automatically create a JMS queue corresponding to the address settings match when a JMS producer or a consumer tries to use such a queue. The default is <b>false</b> . This feature is provided as a <a href="#">technology preview</a> only.
auto-delete-jms-queues	Determines whether JBoss EAP should automatically delete auto-created JMS queues when they have no consumers and no messages. The default is <b>false</b> . This feature is provided as a <a href="#">technology preview</a> only.
dead-letter-address	The address to send dead messages to. See <a href="#">Configuring Dead Letter Addresses</a> for more information.
expiry-address	The address that will receive expired messages. See <a href="#">Configuring Message Expiry</a> for details.
expiry-delay	Defines the expiration time, in milliseconds, that will be used for messages using the default expiration time. Default is <b>-1</b> .
last-value-queue	Defines whether a queue only uses last values or not. See <a href="#">Last-value Queues</a> for more information.
max-delivery-attempts	Defines how many time a canceled message can be redelivered before sending to the dead-letter-address. Default is <b>10</b> .
max-redelivery-delay	Maximum value for the redelivery-delay, in milliseconds. Default is <b>0</b> .
max-size-bytes	The maximum size for this address, in bytes. Default is <b>-1</b> .
message-counter-history-day-limit	Day limit for the message counter history. Default is <b>0</b> .
page-max-cache-size	The number of page files to keep in memory to optimize IO during paging navigation. Default is <b>5</b> .

Name	Description
page-size-bytes	The paging size, in bytes. Default is <b>10485760</b> .
redelivery-delay	Defines how long to wait before attempting redelivery of a canceled message, in milliseconds. Default is <b>0</b> . See <a href="#">Configuring Delayed Redelivery</a> for more information.
redelivery-multiplier	Multiplier to apply to the redelivery-delay parameter. Default is <b>1.0</b> .
redistribution-delay	Defines how long to wait, in milliseconds, after the last consumer is closed on a queue before redistributing any messages. Default is <b>-1</b> .
send-to-dla-on-no-route	When set to <b>true</b> , a message will be sent to the configured dead letter address if it cannot be routed to any queues. Default is <b>false</b> .
slow-consumer-check-period	How often to check, in seconds, for slow consumers. Default is <b>5</b> .
slow-consumer-policy	Determines what happens when a slow consumer is identified. Valid options are <b>KILL</b> or <b>NOTIFY</b> . <b>KILL</b> will kill the consumer's connection, which will impact any client threads using that same connection. <b>NOTIFY</b> will send a <b>CONSUMER_SLOW</b> management notification to the client. Default is <b>NOTIFY</b> .
slow-consumer-threshold	The minimum rate of message consumption allowed before a consumer is considered slow. Default is <b>-1</b> .

## A.2. CONNECTION FACTORY ATTRIBUTES

Table A.2. Connection Factory Attributes

Attribute	Description
auto-group	Whether message grouping is automatically used.
block-on-acknowledge	Whether to block on acknowledge.
block-on-durable-send	Whether to block on durable send.
block-on-non-durable-send	Whether to block on non durable send.
cache-large-message-client	Whether to cache large messages.
call-failover-timeout	The timeout, in milliseconds, to use when failover is in process.
call-timeout	The call timeout, in milliseconds.

Attribute	Description
client-failure-check-period	The client failure check period, in milliseconds.
client-id	The client ID.
compress-large-messages	Whether large messages should be compressed.
confirmation-window-size	The confirmation window size, in bytes.
connection-load-balancing-policy-class-name	Name of a class implementing a client-side load balancing policy that a client can use to load balance sessions across different nodes in a cluster.
connection-ttl	The connection time to live, in milliseconds.
connectors	Defines the connectors, which are stored in a map by connector name (with an undefined value). It is possible to pass a list of connector names when writing this attribute.
consumer-max-rate	The consumer maximum rate, per second.
consumer-window-size	The consumer window size, in bytes.
discovery-group	The discovery group name.
dups-ok-batch-size	The dups ok batch size.
entries	The JNDI names the connection factory should be bound to.
factory-type	<p>The type of connection factory. Valid values are:</p> <ul style="list-style-type: none"> <li>• <b>GENERIC</b></li> <li>• <b>TOPIC</b></li> <li>• <b>QUEUE</b></li> <li>• <b>XA_GENERIC</b></li> <li>• <b>XA_QUEUE</b></li> <li>• <b>XA_TOPIC</b></li> </ul> <p>Use <b>GENERIC</b> for a general connection to a broker, while <b>TOPIC</b> and <b>QUEUE</b> should be used for connections to their respective JMS types. The XA counterparts should be used for transactional messaging.</p>
failover-on-initial-connection	Whether to failover on initial connection.
group-id	The group ID.

Attribute	Description
ha	Whether the connection factory supports high availability.
max-retry-interval	The maximum retry interval, in milliseconds.
min-large-message-size	The minimum large message size, in bytes.
pre-acknowledge	Whether to pre-acknowledge.
producer-max-rate	The producer maximum rate, per second.
producer-window-size	The producer window size, in bytes.
protocol-manager-factory	The protocol manager factory used by this connection factory.
reconnect-attempts	The reconnect attempts.
retry-interval	The retry interval, in milliseconds.
retry-interval-multiplier	The retry interval multiplier.
scheduled-thread-pool-max-size	The scheduled thread pool maximum size.
thread-pool-max-size	The thread pool maximum size.
transaction-batch-size	The transaction batch size.
use-global-pools	Whether to use global pools.

### A.3. POOLED CONNECTION FACTORY ATTRIBUTES

**Table A.3. Pooled Connection Factory Attributes**

Attribute	Description
auto-group	Whether message grouping is automatically used.
block-on-acknowledge	Whether to block on acknowledge.
block-on-durable-send	Whether to block on durable send.
block-on-non-durable-send	Whether to block on non durable send.
cache-large-message-client	Whether to cache large messages.

Attribute	Description
call-failover-timeout	The timeout, in milliseconds, to use when failover is in process.
call-timeout	The call timeout, in milliseconds.
client-failure-check-period	The client failure check period, in milliseconds.
client-id	The client id.
compress-large-messages	Whether large messages should be compressed.
confirmation-window-size	The confirmation window size, in bytes.
connection-load-balancing-policy-class-name	Name of a class implementing a client-side load balancing policy that a client can use to load balance sessions across different nodes in a cluster.
connection-ttl	The connection time to live, in milliseconds.
connectors	Defines the connectors, which are stored in a map by connector name (with an undefined value). It is possible to pass a list of connector names when writing this attribute.
consumer-max-rate	The consumer max rate, per second.
consumer-window-size	The consumer window size, in bytes.
discovery-group	The discovery group name.
dups-ok-batch-size	The dups ok batch size.
enlistment-trace	Enables IronJacamar to record enlistment traces for this pooled connection factory. This attribute is undefined by default and the behavior is driven by the presence of the <b>ironjacamar.disable_enlistment_trace</b> system property.
entries	The JNDI names that the connection factory should be bound to.
failover-on-initial-connection	Whether to failover on initial connection.
group-id	The group id.
ha	Whether the connection factory supports high availability.
initial-connect-attempts	The number of attempts to connect initially with this factory.
initial-message-packet-size	The initial size of messages created through this factory.

Attribute	Description
jndi-params	The JNDI params to use for locating the destination for incoming connections.
managed-connection-pool	The class name of the managed connection pool used by this pooled connection factory.
max-pool-size	The maximum size for the pool.
max-retry-interval	The max retry interval, in milliseconds.
min-large-message-size	The min large message size, in bytes.
min-pool-size	The minimum size for the pool.
password	The default password to use with this connection factory. This is only needed when pointing the connection factory to a remote host.
pre-acknowledge	Whether to pre-acknowledge.
producer-max-rate	The producer max rate, per second.
producer-window-size	The producer window size, in bytes.
protocol-manager-factory	The protocol manager factory used by this pooled connection factory.
reconnect-attempts	The reconnect attempts. By default, a pooled connection factory will try to reconnect infinitely to the messaging servers.
retry-interval	The retry interval, in milliseconds.
retry-interval-multiplier	The retry interval multiplier.
scheduled-thread-pool-max-size	The scheduled thread pool max size.
setup-attempts	The number of times to set up an MDB endpoint.
setup-interval	The interval, in milliseconds, between attempts at setting up an MDB endpoint.
thread-pool-max-size	The thread pool max size.
transaction	The transaction mode.
transaction-batch-size	The transaction batch size.

Attribute	Description
use-auto-recovery	Whether to use auto recovery.
use-global-pools	Whether to use global pools.
use-jndi	Use JNDI to locate the destination for incoming connections.
use-local-tx	Use a local transaction for incoming sessions.
user	The default username to use with this connection factory. This is only needed when pointing the connection factory to a remote host.

## A.4. CORE BRIDGE ATTRIBUTES

Table A.4. Core Bridge Attributes

Attribute	Description
check-period	The period, in milliseconds, between client failure checks.
confirmation-window-size	The size to use for the connection used to forward messages to the target node.
connection-ttl	The maximum time, in milliseconds, for which the connections used by the bridges are considered alive, in the absence of heartbeat.
discovery-group	The name of the discovery group used by this bridge. This attribute may not be set if <b>static-connectors</b> is defined.
filter	An optional filter string. If specified, then only messages that match the filter expression specified will be forwarded.
forwarding-address	The address on the target server that the message will be forwarded to. If a forwarding address is not specified, then the original destination of the message will be retained.
ha	Whether or not this bridge should support high availability. If <b>true</b> , then it will connect to any available server in a cluster and support failover. The default value is <b>false</b> .
initial-connect-attempts	The number of attempts to connect initially with this bridge.
max-retry-interval	The maximum interval of time used to retry connections.
min-large-message-size	The minimum size, in bytes, for a message before it is considered as a large message.



Attribute	Description
password	The password to use when creating the bridge connection to the remote server. If it is not specified, then the default cluster password specified by the <b>cluster-password</b> attribute in the <b>messaging-activemq</b> subsystem resource will be used.
queue-name	The unique name of the local queue that the bridge consumes from. The queue must already exist by the time the bridge is instantiated at startup.
reconnect-attempts	The total number of reconnect attempts that the bridge will make before giving up and shutting down. The default value is <b>-1</b> , which signifies an unlimited number of attempts.
reconnect-attempts-on-same-node	The total number of reconnect attempts on the same node that the bridge will make before giving up and shutting down. A value of <b>-1</b> signifies an unlimited number of attempts. The default is <b>10</b> .
retry-interval	The period, in milliseconds, between subsequent reconnection attempts, if the connection to the target server has failed.
retry-interval-multiplier	A multiplier to apply to the time since the last retry to compute the time to the next retry. This allows you to implement an exponential backoff between retry attempts.
static-connectors	A list of statically defined connectors used by this bridge. This attribute may not be set if <b>discovery-group</b> is defined.
transformer-class-name	The name of a user-defined class that implements the <b>org.apache.activemq.artemis.core.server.cluster.Transformer</b> interface.
use-duplicate-detection	Whether the bridge will automatically insert a duplicate ID property into each message that it forwards.
user	The user name to use when creating the bridge connection to the remote server. If not specified, the default cluster user specified by the <b>cluster-user</b> attribute in the <b>messaging-activemq</b> subsystem resource will be used.

## A.5. JMS BRIDGE ATTRIBUTES

Table A.5. JMS bridge Attributes

Attribute	Description
-----------	-------------

Attribute	Description
add-messageID-in-header	If this is set to <b>true</b> , then the original message's message ID will be appended in the message sent to the destination in the header <b>AMQ_BRIDGE_MSG_ID_LIST</b> . If the message is bridged more than once, each message ID will be appended.
client-id	The JMS client ID to use when creating and looking up the subscription if it is durable and the source destination is a topic.
failure-retry-interval	The amount of time, in milliseconds, to wait between trying to recreate connections to the source or target servers when the bridge has detected they have failed.
max-batch-size	The maximum number of messages to consume from the source destination before sending them in a batch to the target destination. The value must be greater than or equal to <b>1</b> .
max-batch-time	The maximum number of milliseconds to wait before sending a batch of messages to a target, even if the number of messages consumed has not reached max-batch-size. A value of <b>-1</b> means to wait forever.
max-retries	The number of times to attempt to recreate connections to the source or target servers when the bridge has detected they have failed. The bridge will give up after trying this number of times. A value of <b>-1</b> means to try forever.
module	The name of the JBoss EAP module containing the resources required to look up source and target JMS resources.
paused	A read-only property that reports whether the JMS bridge is paused.
quality-of-service	The desired quality of service mode. Possible values are <b>AT_MOST_ONCE</b> , <b>DUPLICATES_OK</b> , or <b>ONCE_AND_ONLY_ONCE</b> . See <a href="#">Quality of Service</a> for details on the different modes.
selector	A JMS selector expression used for consuming messages from the source destination. Only messages that match the selector expression will be bridged from the source to the target destination.
subscription-name	The name of the subscription if it is durable and the source destination is a topic.
source-connection-factory	The name of the source connection factory to look up on the source messaging server.
source-context	The properties used to configure the source JNDI initial context.
source-destination	The name of the source destination to look up on the source messaging server.

Attribute	Description
source-password	The password for creating the source connection.
source-user	The name of the user for creating the source connection.
target-connection-factory	The name of the target connection factory to look up on the target messaging server.
target-context	The properties used to configure the target JNDI initial context.
target-destination	The name of the target destination to look up on the target messaging server.
target-password	The password for creating the target connection.
target-user	The name of the user for creating the target connection.

## A.6. CLUSTER CONNECTION ATTRIBUTES

Table A.6. Cluster Connection Attributes

Attribute	Description
allow-direct-connections-only	If set to <b>true</b> , this node will <b>not</b> create a connection to another node in the cluster if it resides more than 1 hop away. Used only when the attribute <b>static-connectors</b> is defined. The default is <b>false</b> .
call-failover-timeout	The timeout, in milliseconds, to use when failover is in process for remote calls made by the cluster connection. The default is <b>-1</b> , which is unbounded.
call-timeout	The timeout, in milliseconds, for remote calls made by the cluster connection. The default is <b>30000</b> , or 30 seconds.
check-period	The period, in milliseconds, between client failure check. The default is <b>30000</b> , or 30 seconds.
cluster-connection-address	Each cluster connection only applies to messages sent to an address that starts with this value.
confirmation-window-size	The window size, in bytes, for the connection used to forward messages to a target node. The default is <b>1048576</b> .
connection-ttl	The maximum time, in milliseconds, for which the connections used by the cluster connections are considered alive in the absence of heartbeat. The default is <b>60000</b> , or 60 seconds.

Attribute	Description
connector-name	The name of the connector to use for the cluster connection.
discovery-group	The discovery group used to obtain the list of other servers in the cluster to which this cluster connection will make connections. Must be undefined (null) if <b>static-connectors</b> is defined.
initial-connect-attempts	The number of attempts to connect initially with this cluster connection. The default is <b>-1</b> , which is unbounded.
max-hops	The maximum number of times a message can be forwarded. JBoss EAP can be configured to also load balance messages to nodes that might be connected to it only indirectly with other ActiveMQ Artemis messaging servers as intermediates in a chain. The default is <b>1</b> .
max-retry-interval	The maximum interval of time, in milliseconds, used to retry connections. The default is <b>2000</b> , or two seconds.
message-load-balancing-type	<p>This parameter determines how messages will be distributed between other nodes in the cluster. Replaces the deprecated <b>forward-when-no-consumers</b>. Valid values are <b>OFF</b>, <b>STRICT</b>, or <b>ON_DEMAND</b>.</p> <p><b>OFF</b></p> <p>Messages will never be forwarded to another node in the cluster.</p> <p><b>STRICT</b></p> <p>Messages will be distributed in a round robin fashion even though the same queues on the other nodes of the cluster may have no consumers at all, or they may have consumers that have nonmatching message filters or selectors. Note that JBoss EAP will not forward messages to other nodes if there are no queues of the same name on the other nodes, even if this parameter is set to <b>STRICT</b>. Using <b>STRICT</b> is like setting the legacy <b>forward-when-no-consumers</b> parameter to <b>true</b>.</p> <p><b>ON_DEMAND</b></p> <p>Messages are forwarded to other nodes of the cluster if the forwarding address has queues that have consumers. If those consumers have message filters or selectors, at least one of those selectors must match the message. Using <b>ON_DEMAND</b> is like setting the legacy <b>forward-when-no-consumers</b> parameter to <b>false</b>.</p> <p>The default is <b>ON_DEMAND</b>.</p>
min-large-message-size	The minimum size, in bytes, for a message before it is considered as a large message. The default is <b>102400</b> .
node-id	The node ID used by this cluster connection. This attribute is read only.
notification-attempts	How many times the cluster connection will broadcast itself. The default is <b>2</b> .

Attribute	Description
notification-interval	The interval, in milliseconds, between notifications. The default is <b>10000</b> , or 10 seconds.
reconnect-attempts	The total number of reconnect attempts the bridge will make before giving up and shutting down. The default is <b>-1</b> , which signifies an unlimited number of attempts.
retry-interval	The period, in milliseconds, between subsequent attempts to reconnect to a target server, if the connection to the target server has failed. The default is <b>500</b> .
retry-interval-multiplier	A multiplier to apply to the time since the last retry to compute the time to the next retry. This allows you to implement an exponential backoff between retry attempts. The default is <b>1.0</b> .
static-connectors	The statically defined list of connectors to which this cluster connection will make connections. Must be undefined if <b>discovery-group-name</b> is defined.
topology	The topology of the nodes that this cluster connection is aware of. This attribute is read only.
use-duplicate-detection	Whether the bridge will automatically insert a duplicate ID property into each message that it forwards. The default is <b>true</b> .

*Revised on 2018-02-08 10:16:22 EST*