

The logo for FuseSource, featuring the word "FuseSource" in a white, sans-serif font, oriented vertically on a black background. The background has a jagged, torn-paper edge effect with red and yellow colors visible underneath.

FuseSource

Apache ActiveMQ
Fault Tolerant Messaging

Version 7.1
December 2012

Integration Everywhere

Fault Tolerant Messaging

Version 7.1

Updated: 08 Jan 2014

Copyright © 2012 Red Hat, Inc. and/or its affiliates.

Trademark Disclaimer

These materials and all Progress® software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

Fuse, FuseSource, Fuse ESB, Fuse ESB Enterprise, Fuse MQ Enterprise, Fuse Mediation Router, Fuse Message Broker, Fuse Services Framework, Fuse IDE, Fuse HQ, Fuse Management Console, and Integration Everywhere are trademarks or registered trademarks of FuseSource Corp. or its parent corporation, Progress Software Corporation, or one of their subsidiaries or affiliates in the United States. Apache, ServiceMix, Camel, CXF, and ActiveMQ are trademarks of Apache Software Foundation. Any other names contained herein may be trademarks of their respective owners.

Third Party Acknowledgements

One or more products in the Fuse ESB Enterprise release includes third party components covered by licenses that require that the following documentation notices be provided:

- JLine (<http://jline.sourceforge.net>) jline:jline:jar:1.0

License: BSD (LICENSE.txt) - Copyright (c) 2002-2006, Marc Prud'hommeaux <mwpl@cornell.edu>

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of JLine nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Stax2 API (<http://woodstox.codehaus.org/StAX2>) org.codehaus.woodstox:stax2-api:jar:3.1.1

License: The BSD License (<http://www.opensource.org/licenses/bsd-license.php>)

Copyright (c) <YEAR>, <OWNER> All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- jibx-run - JiBX runtime (<http://www.jibx.org/main-reactor/jibx-run>) org.jibx:jibx-run:bundle:1.2.3

License: BSD (<http://jibx.sourceforge.net/jibx-license.html>) Copyright (c) 2003-2010, Dennis M. Sosnoski.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of JiBX nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- JavaAssist (<http://www.jboss.org/javassist>) org.jboss.javassist:com.springsource.javassist:jar:3.9.0.GA:compile

License: MPL (<http://www.mozilla.org/MPL/MPL-1.1.html>)

- HAPI-OSGI-Base Module (<http://hl7api.sourceforge.net/hapi-osgi-base/>) ca.uhn.hapi:hapi-osgi-base:bundle:1.2
License: Mozilla Public License 1.1 (<http://www.mozilla.org/MPL/MPL-1.1.txt>)

Table of Contents

1. Introduction	11
2. Client Failover	13
Failover Protocol	14
Static Failover	15
Dynamic Failover	18
Discovery Protocol	22
Discovery URI	23
Discovery Agents	25
Fuse Fabric Discovery Agent	27
Static Discovery Agent	28
Multicast Discovery Agent	29
Zeroconf Discovery Agent	31
3. Master/Slave	33
Shared File System Master/Slave	34
Shared JDBC Master/Slave	39
4. Master/Slave and Broker Networks	45
Index	49

List of Figures

3.1. Shared File System Initial State	35
3.2. Shared File System after Master Failure	36
3.3. Shared File System after Master Restart	38
3.4. JDBC Master/Slave Initial State	40
3.5. JDBC Master/Slave after Master Failure	41
3.6. JDBC Master/Slave after Master Restart	43
4.1. Master/Slave Groups on Two Host Machines	46
4.2. Broker Network Consisting of Host Pairs	47

List of Tables

2.1. Failover Transport Options	15
2.2. Broker-side Failover Properties	19
2.3. Dynamic Discovery Protocol Options	23

List of Examples

2.1. Simple Failover URI	17
2.2. Broker for Dynamic Failover	20
2.3. Failover URI for Connecting to a Failover Cluster	21
2.4. Discovery URI	23
2.5. Discovery Protocol URI	24
2.6. Injecting Transport Options into a Discovered Transport	24
2.7. Enabling a Discovery Agent on a Broker	25
2.8. Fuse Fabric Discovery Agent URI Format	27
2.9. Client Connection URL using Fuse Fabric Discovery	27
2.10. Static Discovery Agent URI Format	28
2.11. Discovery URI using the Static Discovery Agent	28
2.12. Multicast Discovery Agent URI Format	29
2.13. Enabling a Multicast Discovery Agent on a Broker	29
2.14. Client Connection URL using Multicast Discovery	30
2.15. Zeroconf Discovery Agent URI Format	31
2.16. Enabling a Multicast Discovery Agent on a Broker	32
2.17. Client Connection URL using Zeroconf Discovery	32
3.1. Shared File System Broker Configuration	37
3.2. Client URL for a Shared File System Master/Slave Group	37
3.3. JDBC Master/Slave Broker Configuration	42
3.4. Client URL for a Shared JDBC Master/Slave Group	43
4.1. Network Connector to a Master/Slave Group	46

Chapter 1. Introduction

Fault tolerant message systems can recover from failures with little or no interruption of functionality. Fuse MQ Enterprise does this by making it easy to configure clients to fail over to new brokers in the event of a broker failure. It also makes it easy to set up master/slave groups that allow brokers to take over for each other and maintain the integrity of persistent messages and transactions.

Overview

If planned for, disaster scenarios that result in the loss of a message broker need not obstruct message delivery. Making a messaging system fault tolerant involves:

- deploying multiple brokers into a topology that allows one broker to pick up the duties of a failed broker
- configuring clients to fail over to a new broker in the event that its current broker fails

Fuse MQ Enterprise provides mechanisms that make building fault tolerant messaging systems easy.

Client fail over

Fuse MQ Enterprise provides two protocols that allow clients to fail over to a new broker in the case of a failure:

- the failover protocol—allows you to provide a list of brokers that a client can use
- the discovery protocol—allows clients to automatically discover the brokers available for fail over

Both protocols automatically reconnect to an available broker when its existing connection fails. As long as an available broker is running, the client can continue to function uninterrupted.

When combined with brokers deployed in a master/slave topology, the failover protocol is a key part of a fault-tolerant messaging system. The clients will automatically fail over to the slave broker if the master fails. The clients will remain functional and continue working as if nothing had happened.

For more information, see ["Client Failover" on page 13](#).

Master/Slave topologies

A *master/slave* topology includes a master broker and one or more slave brokers. All of the brokers share data by using either a replication mechanism or by using a shared data store. When the master broker fails, one of the slave brokers takes over and becomes the new master broker. Client applications can reconnect to the new master broker and resume processing as normal.

For details, see ["Master/Slave" on page 33](#).

Chapter 2. Client Failover

Fuse MQ Enterprise provides two simple mechanisms for clients to failover to an alternate broker if its active connection fails. The failover protocol relies on a hard coded list of alternative brokers. The discovery protocol relies on discovery agents to provide a list of alternative brokers.

Failover Protocol	14
Static Failover	15
Dynamic Failover	18
Discovery Protocol	22
Discovery URI	23
Discovery Agents	25
Fuse Fabric Discovery Agent	27
Static Discovery Agent	28
Multicast Discovery Agent	29
Zeroconf Discovery Agent	31

Failover Protocol

Static Failover	15
Dynamic Failover	18

The *failover protocol* facilitates quick recovery from network failures. When a recoverable network error occurs the protocol catches the error and automatically attempts to reestablish the connection to an alternate broker endpoint without the need to recreate all of the objects associated with the connection. The failover URI is composed of one or more URIs that represent different broker endpoints. By default, the protocol randomly chooses a URI from the list and attempts to establish a network connection to it. If it does not succeed, or if it subsequently fails, a new network connection is established to one of the other URIs in the list.

You can set up failover in one of the following ways:

- **Static**—the client is configured with a static list of available URIs
- **Dynamic**—the brokers push information about the available broker connections

Static Failover

Overview

In static failover a client is configured to use a *failover URI* that lists the URIs of the broker connections the client can use. When establishing a connection, the client randomly chooses a URI from the list and attempts to establish a connection to it. If the connection does not succeed, the client chooses a new URI from the list and tries again. The client will continue cycling through the list until a connection attempt succeeds.

If a client's connection to a broker fails after it has been established, the client will attempt to reconnect to a different broker in the list. Once a connection to a new broker is established, the client will continue to use the new broker until the connection to the new broker is severed.

Failover URI

A failover URI is a composite URI that uses one of the following syntaxes:

- `failover:uri1,...,uriN`
- `failover:(uri1,...,uriN)?TransportOptions`

The URI list(*uri1, ..., uriN*) is a comma-separated list containing the list of broker endpoint URIs to which the client can connect. The transport options(*?TransportOptions*) specified in the form of a query list, allow you to configure some of the failover behaviors.

Transport options

The failover protocol supports the transport options described in [Table 2.1 on page 15](#).

Table 2.1. Failover Transport Options

Option	Default	Description
<code>initialReconnectDelay</code>	10	Specifies the number of milliseconds to wait before the first reconnect attempt.
<code>maxReconnectDelay</code>	30000	Specifies the maximum amount of time, in milliseconds, to wait between reconnect attempts.

Option	Default	Description
<code>useExponentialBackOff</code>	<code>true</code>	Specifies whether to use an exponential back-off between reconnect attempts.
<code>backOffMultiplier</code>	<code>2</code>	Specifies the exponent used in the exponential back-off algorithm.
<code>maxReconnectAttempts</code>	<code>-1</code>	Specifies the maximum number of reconnect attempts before an error is returned to the client. <code>-1</code> specifies unlimited attempts. <code>0</code> specifies that an initial connection attempt is made at start-up, but no attempts to failover over to a secondary broker will be made.
<code>startupMaxReconnectAttempts</code>	<code>0</code>	Specifies the maximum number of reconnect attempts before an error is returned to the client on the <i>first</i> attempt by the client to start a connection. <code>0</code> specifies unlimited attempts.
<code>randomize</code>	<code>true</code>	Specifies if a URI is chosen at random from the list. Otherwise, the list is traversed from left to right.
<code>backup</code>	<code>false</code>	Specifies if the protocol initializes and holds a second transport connection to enable fast failover.
<code>timeout</code>	<code>-1</code>	Specifies the amount of time, in milliseconds, to wait before sending an error if a new connection is not established. <code>-1</code> specifies an infinite timeout value.
<code>trackMessages</code>	<code>false</code>	Specifies if the protocol keeps a cache of in-flight messages that are flushed to a broker on reconnect.
<code>maxCacheSize</code>	<code>131072</code>	Specifies the size, in bytes, used for the cache used to track messages.
<code>updateURIsSupported</code>	<code>true</code>	Specifies whether the client accepts updates to its list of known URIs from the connected broker. Setting this to

Option	Default	Description
		false inhibits the client's ability to use dynamic failover. See "Dynamic Failover" on page 18 .
updateURIsURL		Specifies a URL locating a text file that contains a comma-separated list of URIs to use for reconnect in the case of failure. See "Dynamic Failover" on page 18 .

Example

[Example 2.1 on page 17](#) shows a failover URI that can connect to one of two message brokers.

Example 2.1. Simple Failover URI

```
failover:(tcp://localhost:61616,tcp://remotehost:61616)?initialReconnectDelay=100
```

Dynamic Failover

Overview

Dynamic failover combines the failover protocol and a network of brokers to allow a broker to supply its clients with a list of broker connections to which the clients can failover. Clients use a failover URI to connect to a broker and the broker dynamically updates the clients' list of available URIs. The broker updates its clients' failover lists with the URIs of the other brokers in its network of brokers that are currently running. As new brokers join, or exit, the network of brokers, the broker will adjust its clients' failover lists.

From a connectivity point of view, dynamic failover works the same as static failover. A client randomly chooses a URI from the list provided in its failover URI. Once that connection is established, the list of available brokers is updated. If the original connection fails, the client will randomly select a new URI from its dynamically generated list of brokers. If the new broker is configured for to supply a failover list, the new broker will update the client's list.

Set-up

To use dynamic failover you must configure both the clients and brokers used by your application. The following must be configured:

- The client's must be configured to use the failover protocol when connecting with its broker.
 - The client must be configured to accept URI lists from a broker.
 - The brokers must be configured to form a network of brokers.

See [Using Networks of Brokers](#).
 - The broker's transport connector must set the failover properties needed to update its consumers.
-

Client-side configuration

The client-side configuration for using dynamic failover is nearly identical to the client-side configuration for using static failover. The differences include:

- The failover URI can consist of a single broker URI.
- The `updateURIsSupported` option must be set to `true`.

- The `updateURIsURL` option should be set so that the transport can failover to a new broker when none of the broker's in the dynamically supplied list are available.

See ["Failover URI" on page 15](#) and ["Transport options" on page 15](#) for more information about using failover URIs.

Broker-side configuration



Important

Brokers should *never* use a failover URI to configure a transport connector. The failover protocol does not support listening for incoming messages.

Configuring a broker to participate in dynamic failover requires two things:

- The broker must be configured to participate in a network of brokers that can be available for failovers.

See [Using Networks of Brokers](#) for information about setting up a network of brokers.

- The broker's transport connector must set the failover properties needed to update its consumers.

[Table 2.2 on page 19](#) describes the broker-side properties that can be used to configure a failover cluster. These properties are attributes on the broker's `transportConnector` element.

Table 2.2. Broker-side Failover Properties

Property	Default	Description
<code>updateClusterClients</code>	<code>false</code>	Specifies if the broker passes information to connected clients about changes in the topology of the broker cluster.
<code>updateClusterClientsOnRemove</code>	<code>false</code>	Specifies if clients are updated when a broker is removed from the cluster.
<code>rebalanceClusterClients</code>	<code>false</code>	Specifies if connected clients are asked to rebalance across the cluster whenever a new broker joins.
<code>updateClusterFilter</code>		Specifies a comma-separated list of regular expression filters, which

Property	Default	Description
		match against broker names to select the brokers that belong to the failover cluster.

Example

[Example 2.2 on page 20](#) shows the configuration for a broker that participates in dynamic failover.

Example 2.2. Broker for Dynamic Failover

```

<beans ... >
  <broker>
    ...
    <networkConnectors>
      ❶ <networkConnector uri="multicast://default" />
    </networkConnectors>
    ...
    <transportConnectors>
      <transportConnector name="openwire"
        uri="tcp://0.0.0.0:61616"
        ❷ discoveryUri="multicast://default"
        ❸ updateClusterClients="true"
        ❹ updateClusterFilter="*A*,*B*" />
    </transportConnectors>
    ...
  </broker>
</beans>

```

The configuration in [Example 2.2 on page 20](#) does the following:

- ❶ Creates a network connector that connects to any discoverable broker that uses the multicast transport.
- ❷ Makes the broker discoverable by other brokers over the multicast protocol.
- ❸ Makes the broker update the list of available brokers for clients that connect using the failover protocol.

**Note**

Clients will only be updated when new brokers join the cluster, not when a broker leaves the cluster.

- ❶ Creates a filter so that only those brokers whose names start with the letter `A` or the letter `B` are considered to belong to the failover cluster.

[Example 2.3 on page 21](#) shows the URI for a client that uses the failover protocol to connect to the broker and its cluster.

Example 2.3. Failover URI for Connecting to a Failover Cluster

```
failover:(tcp://0.0.0.0:61616)?initialReconnectDelay=100
```

Discovery Protocol

Discovery URI	23
Discovery Agents	25
Fuse Fabric Discovery Agent	27
Static Discovery Agent	28
Multicast Discovery Agent	29
Zeroconf Discovery Agent	31

The failover protocol provides a lot of control over the brokers to which a client can connect. Using dynamic failover adds some ability to make the broker list more transparent. However, it has weaknesses. It requires that you know the address of at least one broker and that an initial broker is active when the client starts up. Using dynamic failover also requires that all of the brokers being used for failover are configured in a network of brokers.

Fuse MQ Enterprise's discovery protocol offers an alternative method for dynamically generating a list of brokers that are available for client failover. The protocol allows brokers to advertise their availability and for clients to dynamically discover them. This is accomplished using two pieces:

- *discovery URI*—looks up all of the discoverable brokers and presents them as a list of actual URIs for use by the client or network connector
- *discovery agents*—components that advertise the list of available brokers

Discovery URI

Overview

The discovery URI is a virtual URI that specifies which discovery agent to use for discovering available brokers. The discovery protocol connects to the specified agent and uses that data returned from the agent to build up a list of broker URIs.

URI syntax

[Example 2.4 on page 23](#) shows the syntax for a discovery URI.

Example 2.4. Discovery URI

```
discovery://(DiscoveryAgentUri)?Options
```

DiscoveryAgentUri is URI for the discovery agent used to build up the list of available brokers. Discovery agents are described in ["Discovery Agents" on page 25](#).

The options, *?Options*, are specified in the form of a query list. The discovery options are described in [Table 2.3 on page 23](#). You can also inject transport options as described in ["Setting options on the discovered transports" on page 24](#).



Tip

If no options are required, you can drop the parentheses from the URI. The resulting URI would take the form

```
discovery://DiscoveryAgentUri
```

Transport options

The discovery protocol supports the options described in [Table 2.3 on page 23](#).

Table 2.3. Dynamic Discovery Protocol Options

Option	Default	Description
<code>initialReconnectDelay</code>	10	Specifies, in milliseconds, how long to wait before the first reconnect attempt.
<code>maxReconnectDelay</code>	30000	Specifies, in milliseconds, the maximum amount of time to wait between reconnect attempts.

Option	Default	Description
<code>useExponentialBackOff</code>	<code>true</code>	Specifies if an exponential back-off is used between reconnect attempts.
<code>backOffMultiplier</code>	2	Specifies the exponent used in the exponential back-off algorithm.
<code>maxReconnectAttempts</code>	0	Specifies the maximum number of reconnect attempts before an error is sent back to the client. 0 specifies unlimited attempts.

Sample URI

[Example 2.5 on page 24](#) shows a discovery URI that uses a multicast discovery agent.

Example 2.5. Discovery Protocol URI

```
discovery://(multicast://default)?initialReconnectDelay=100
```

Setting options on the discovered transports

The list of transport options, *Options*, in the discovery URI can also be used to set options on the *discovered* transports. If you set an option *not* listed in [Table 2.3 on page 23](#), the URI parser attempts to inject the option setting into every one of the discovered endpoints.

[Example 2.6 on page 24](#) shows a discovery URI that sets the TCP `connectionTimeout` option to 10 seconds.

Example 2.6. Injecting Transport Options into a Discovered Transport

```
discovery://(multicast://default)?connectionTimeout=10000
```

The 10 second timeout setting is injected into every discovered TCP endpoint.

Discovery Agents

Fuse Fabric Discovery Agent	27
Static Discovery Agent	28
Multicast Discovery Agent	29
Zeroconf Discovery Agent	31

A discovery agent is a mechanism that advertises available brokers to clients and other brokers. When a client, or broker, using a discovery URI starts up it will look for any brokers that are available using the specified discovery agent. The clients will update their lists periodically using the same mechanism.

How a discovery agent learns about the available brokers varies between agents. Some agents use a static list, some use a third party registry, and some rely on the brokers to provide the information. For discovery agents that rely on the brokers for information, it is necessary to enable the discovery agent in the message broker configuration. For example, to enable the multicast discovery agent on an Openwire endpoint, you edit the relevant `transportConnector` element as shown in [Example 2.7 on page 25](#).

Example 2.7. Enabling a Discovery Agent on a Broker

```
<transportConnectors>
  <transportConnector name="openwire"
    uri="tcp://localhost:61716"
    discoveryUri="multicast://default" />
</transportConnectors>
```

Where the `discoveryUri` attribute on the `transportConnector` element is initialized to `multicast://default`.



Tip

If a broker uses multiple transport connectors, you need to configure each transport connector to use a discovery agent individually. This means that different connectors can use different discovery mechanisms or that one or more of the connectors can be indistinguishable.

Fuse MQ Enterprise currently supports the following discovery agents:

- [Fuse Fabric Discovery Agent](#)
- [Static Discovery Agent](#)

- [Multicast Discovery Agent](#)
- [Zeroconf Discovery Agent](#)

Fuse Fabric Discovery Agent

Overview

The *Fuse Fabric discovery agent* uses Fuse Fabric to discover the brokers in a specified group. The discovery agent requires that all of the discoverable brokers be deployed into a single fabric. When the client attempts to connect to a broker the agent looks up all of the available brokers in the fabric's registry and returns the ones in the specified group.

URI

The Fuse Fabric discovery agent URI conforms to the syntax in [Example 2.8 on page 27](#).

Example 2.8. Fuse Fabric Discovery Agent URI Format

```
fabric://GID
```

Where *GID* is the ID of the broker group from which the client discovers the available brokers.

Configuring a broker

The Fuse Fabric discovery agent requires that the discoverable brokers are deployed into a single fabric.

The best way to deploy brokers into a fabric is using Fuse Management Console. For information on using Fuse Management Console see [Fuse Management Console Documentation](#)¹.

You can also use the console to deploy brokers into a fabric. See "[Fabric Console Commands](#)" in *Console Reference*.

Configuring a client

To use the agent a client must be configured to connect to a broker using a discovery protocol that uses a Fuse Fabric agent URI as shown in [Example 2.9 on page 27](#).

Example 2.9. Client Connection URL using Fuse Fabric Discovery

```
discovery://(fabric://nwBrokers)
```

A client using the URL in [Example 2.9 on page 27](#) will discover all the brokers in the `nwBrokers` broker group and generate a list of brokers to which it can connect.

¹ <http://fusesource.com/docs/fmc>

Static Discovery Agent

Overview

The *static discovery agent* does not truly discover the available brokers. It uses an explicit list of broker URLs to specify the available brokers. Brokers are not involved with the static discovery agent. The client only knows about the brokers that are hard coded into the agent's URI.

Using the agent

The static discovery agent is a client-side only agent. It does not require any configuration on the brokers that will be discovered.

To use the agent, you simply configure the client to connect to a broker using a discovery protocol that uses a static agent URI.

The static discovery agent URI conforms to the syntax in [Example 2.10 on page 28](#).

Example 2.10. Static Discovery Agent URI Format

```
static://(URI1,URI2,URI3,...)
```

Example

[Example 2.11 on page 28](#) shows a discovery URI that configures a client to use the static discovery agent to connect to one member of a broker pair.

Example 2.11. Discovery URI using the Static Discovery Agent

```
discovery://(static://(tcp://localhost:61716,tcp://local  
host:61816))
```

Multicast Discovery Agent

Overview

The *multicast discovery agent* uses the IP multicast protocol to find any message brokers currently active on the local network. The agent requires that *each* broker you want to advertise is configured to use the multicast agent to publish its details to a multicast group. Clients using the multicast agent as part of the discovery URI they use for connecting to a broker will use the agent to receive the list of available brokers advertising in the specified multicast group.



Important

Your local network (LAN) must be configured appropriately for the IP/multicast protocol to work.

URI

The multicast discovery agent URI conforms to the syntax in [Example 2.12 on page 29](#).

Example 2.12. Multicast Discovery Agent URI Format

```
multicast://GroupID
```

Where *GroupID* is an alphanumeric identifier. All participants in the same discovery group must use the same *GroupID*.

Configuring a broker

For a broker to be discoverable using the multicast discovery agent, you must enable the discovery agent in the broker's configuration. To enable the multicast discovery agent you set the `transportConnector` element's `discoveryUri` attribute to a multicast discovery agent URI as shown in [Example 2.13 on page 29](#).

Example 2.13. Enabling a Multicast Discovery Agent on a Broker

```
<transportConnectors>
  <transportConnector name="openwire"
    uri="tcp://localhost:61716"
    discoveryUri="multicast://default" />
</transportConnectors>
```

The broker configured in [Example 2.13 on page 29](#) is discoverable as part of the multicast group `default`.

Configuring a client

To use the multicast agent a client must be configured to connect to a broker using a discovery URI that uses a multicast agent URI as shown in [Example 2.14 on page 30](#).

Example 2.14. Client Connection URL using Multicast Discovery

```
discovery://(multicast://default)
```

A client using the URI in [Example 2.14 on page 30](#) will discover all the brokers advertised in the `default` multicast group and generate a list of brokers to which it can connect.

Zeroconf Discovery Agent

Overview

The *zeroconf discovery agent* is derived from Apple's [Bonjour Networking](http://developer.apple.com/networking/bonjour/)² technology, which defines the zeroconf protocol as a mechanism for discovering services on a network. Fuse MQ Enterprise bases its implementation of the zeroconf discovery agent on [JmDNS](http://sourceforge.net/projects/jmdns/)³, which is a service discovery protocol that is layered over IP/multicast and is compatible with Apple Bonjour.

The agent requires that *each* broker you want to advertise is configured to use a multicast discovery agent to publish its details to a multicast group. Clients using the zeroconf agent as part of the discovery URI they use for connecting to a broker will use the agent to receive the list of available brokers advertising in the specified multicast group.



Important

Your local network (LAN) must be configured to use IP/multicast for the zeroconf agent to work.

URI

The zeroconf discovery agent URI conforms to the syntax in [Example 2.15 on page 31](#).

Example 2.15. Zeroconf Discovery Agent URI Format

```
zeroconf://GroupID
```

Where the *GroupID* is an alphanumeric identifier. All participants in the same discovery group must use the same *GroupID*.

Configuring a broker

For a broker to be discoverable using the zeroconf discovery agent, you must enable a multicast discovery agent in the broker's configuration. To enable the multicast discovery agent you set the `transportConnector` element's `discoveryUri` attribute to a multicast discovery agent URI as shown in [Example 2.16 on page 32](#).

² <http://developer.apple.com/networking/bonjour/>

³ <http://sourceforge.net/projects/jmdns/>

Example 2.16. Enabling a Multicast Discovery Agent on a Broker

```
<transportConnectors>
  <transportConnector name="openwire"
    uri="tcp://localhost:61716"
    discoveryUri="multicast://NEGroup" />
</transportConnectors>
```

The broker configured in [Example 2.16 on page 32](#) is discoverable as part of the multicast group `NEGroup`.

Configuring a client

To use the agent a client must be configured to connect to a broker using a discovery protocol that uses a zeroconf agent URI as shown in [Example 2.17 on page 32](#).

Example 2.17. Client Connection URL using Zeroconf Discovery

```
discovery://(zeroconf://NEGroup)
```

A client using the URL in [Example 2.17 on page 32](#) will discover all the brokers advertised in the `NEGroup` multicast group and generate a list of brokers to which it can connect.

Chapter 3. Master/Slave

Persistent messages require an additional layer of fault tolerance. In case of a broker failure, persistent messages require that the replacement broker has a copy of all the undelivered messages. Master/slave groups address this requirement by having a standby broker that shares the active broker's data store.

Shared File System Master/Slave	34
Shared JDBC Master/Slave	39

A master/slave group consists of two or more brokers where one master broker is active and one or more slave brokers are on hot standby, ready to take over whenever the master fails or shuts down. All of the brokers store the message and event data processed by the master broker. So, when one of the slaves takes over as the new master broker the integrity of the messaging system is guaranteed.

Fuse MQ Enterprise supports two master/slave broker configurations:

- [Shared file system](#)—the master and the slaves use a common persistence store that is located on a shared file system
- [Shared JDBC database](#)—the masters and the slaves use a common JDBC persistence store

Shared File System Master/Slave

Overview

A shared file system master/slave group works by sharing a common data store that is located on a shared file system. Brokers automatically configure themselves to operate in master mode or slave mode based on their ability to grab an exclusive lock on the underlying data store.

The disadvantage of this configuration is that the shared file system is a single point of failure. This disadvantage can be mitigated by using a storage area network(SAN) with built in high availability(HA) functionality. The SAN will handle replication and fail over of the data store.

File locking requirements

The shared file system requires an efficient and reliable file locking mechanism to function correctly. Not all SAN file systems are compatible with the shared file system configuration's needs.



Warning

OCFS2 is incompatible with this master/slave configuration, because mutex file locking from Java is not supported.



Warning

NFSv3 is incompatible with this master/slave configuration. In the event of an abnormal termination of a master broker, which is an NFSv3 client, the NFSv3 server does not time out the lock held by the client. This renders the Fuse MQ Enterprise data directory inaccessible. Because of this, the slave broker cannot acquire the lock and therefore cannot start up. In this case, the only way to unblock the master/slave in NFSv3 is to reboot all broker instances.


On the other hand, NFSv4 *is* compatible with this master/slave configuration, because its design includes timeouts for locks. When an NFSv4 client holding a lock terminates abnormally, the lock is automatically released after 30 seconds, allowing another NFSv4 client to grab the lock.

Initial state

[Figure 3.1 on page 35](#) shows the initial state of a shared file system master/slave group. When all of the brokers are started, one of them grabs the exclusive lock on the broker data store and becomes the master. All of the other brokers remain slaves and pause while waiting for the exclusive

lock to be freed up. Only the master starts its transport connectors, so all of the clients connect to it.

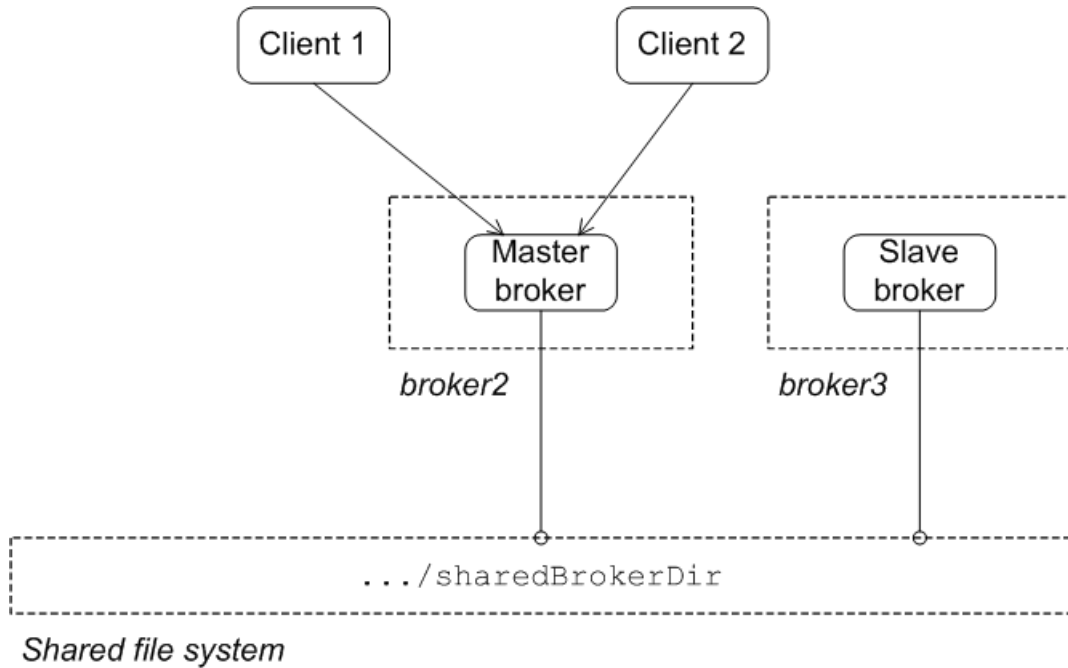
Figure 3.1. Shared File System Initial State



State after failure of the master

Figure 3.2 on page 36 shows the state of the master/slave group after the original master has shut down or failed. As soon as the master gives up the lock (or after a suitable timeout, if the master crashes), the lock on the data store frees up and another broker grabs the lock and gets promoted to master.

Figure 3.2. Shared File System after Master Failure



After the clients lose their connection to the original master, they automatically try all of the other brokers listed in the failover URL. This enables them to find and connect to the new master.

Configuring the brokers

In the shared file system master/slave configuration, there is nothing special to distinguish a master broker from the slave brokers. The membership of a particular master/slave group is defined by the fact that all of the brokers in the group use the *same* persistence layer and store their data in the *same* shared directory.

[Example 3.1 on page 37](#) shows the broker configuration for a shared file system master/slave group that shares a data store located at `/sharedFileSystem/sharedBrokerData` and uses the KahaDB persistence store.

Example 3.1. Shared File System Broker Configuration

```
<broker ... >
...
<persistenceAdapter>
  <kahaDB directory="/sharedFileSystem/sharedBrokerData"/>
</persistenceAdapter>
...
</broker>
```

All of the brokers in the group *must* share the same `persistenceAdapter` element.

Configuring the clients

Clients of shared file system master/slave group must be configured with a failover URL that lists the URLs for all of the brokers in the group.

[Example 3.2 on page 37](#) shows the client failover URL for a group that consists of three brokers: `broker1`, `broker2`, and `broker3`.

Example 3.2. Client URL for a Shared File System Master/Slave Group

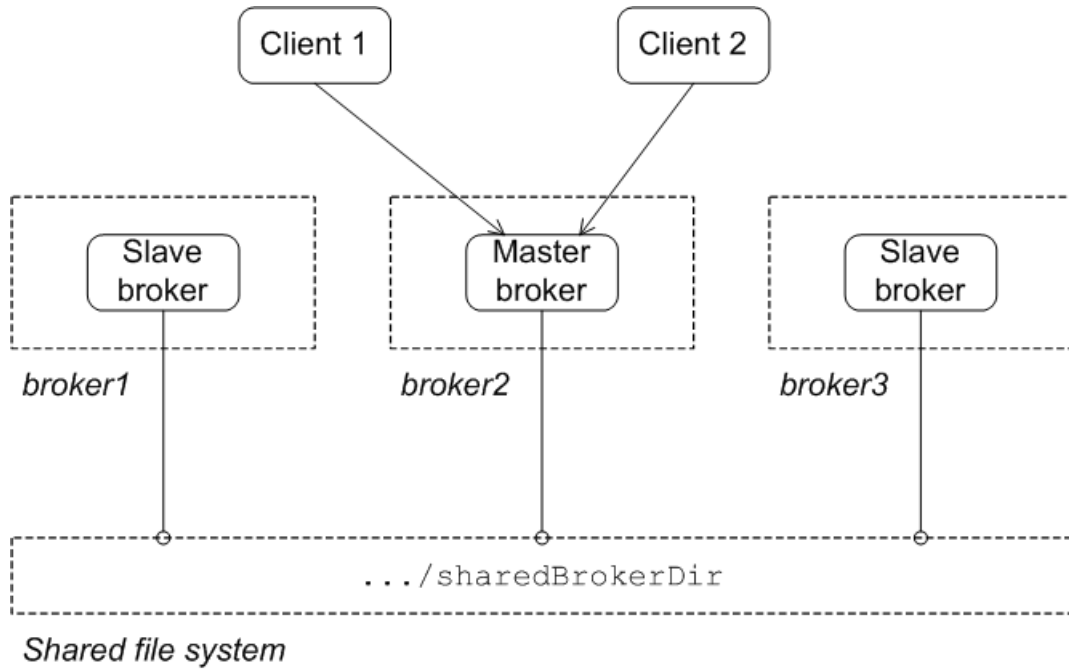
```
fail
over: (tcp://broker1:61616,tcp://broker2:61616,tcp://broker3:61616)
```

For more information about using the failover protocol see ["Static Failover" on page 15](#).

Reintroducing a failed node

You can restart the failed master at any time and it will rejoin the cluster. It will rejoin as a slave broker because one of the other brokers already owns the exclusive lock on the data store, as shown in [Figure 3.3 on page 38](#).

Figure 3.3. Shared File System after Master Restart



Shared JDBC Master/Slave

Overview

A shared JDBC master/slave group works by sharing a common database using the JDBC persistence adapter. Brokers automatically configure themselves to operate in master mode or slave mode, depending on whether or not they manage to grab a mutex lock on the underlying database table.


The disadvantages of this configuration are:

- The shared database is a single point of failure. This disadvantage can be mitigated by using a database with built in high availability(HA) functionality. The database will handle replication and fail over of the data store.
- You cannot enable high speed journaling. This has a significant impact on performance.

Initial state

[Figure 3.4 on page 40](#) shows the initial state of a JDBC master/slave group. When all of the brokers are started, one of them grabs the mutex lock on the database table and becomes the master. All of the other brokers become slaves and pause while waiting for the lock to be freed up. Only the master starts its transport connectors, so all of the clients connect to it.

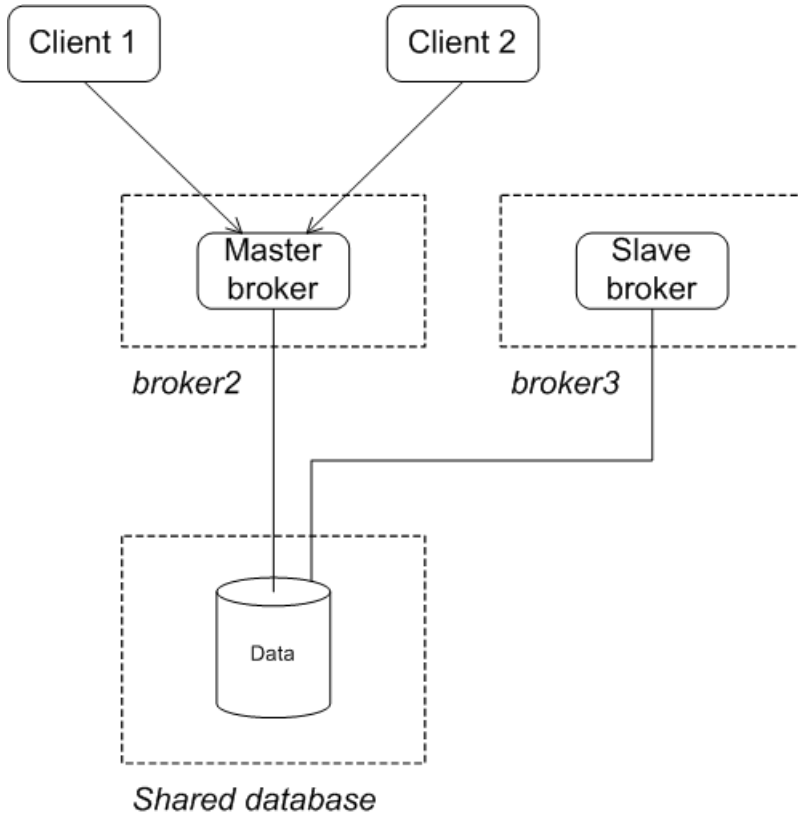
Figure 3.4. JDBC Master/Slave Initial State



After failure of the master

Figure 3.5 on page 41 shows the state of the group after the original master has shut down or failed. As soon as the master gives up the lock (or after a suitable timeout, if the master crashes), the lock on the database table frees up and another broker grabs the lock and gets promoted to master.

Figure 3.5. JDBC Master/Slave after Master Failure



After the clients lose their connection to the original master, they automatically try all of the other brokers listed in the failover URL. This enables them to find and connect to the new master.

Configuring the brokers

In a JDBC master/slave configuration, there is nothing special to distinguish a master broker from the slave brokers. The membership of a particular master/slave group is defined by the fact that all of the brokers in the cluster use the *same* JDBC persistence layer and store their data in the *same* database tables.

There is one important requirement when configuring the JDBC persistence adapter for use in a shared database master/slave cluster. You **must** use the

direct JDBC persistence adapter. This is because the journal used by the journaled JDBC persistence adapter is not replicated and batch updates are used to sync with the JDBC store. Therefore it is not possible to guarantee that the latest updates are on the shared JDBC store.

[Example 3.3 on page 42](#) shows the configuration used by a master/slave group that stores the shared broker data in an Oracle database.

Example 3.3. JDBC Master/Slave Broker Configuration

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:amq="http://activemq.apache.org/schema/core"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
    http://activemq.apache.org/schema/core
    http://activemq.apache.org/schema/core/activemq-core-5.3.1.xsd">

  <broker xmlns="http://activemq.apache.org/schema/core"
    brokerName="brokerA">
    ...
    <persistenceAdapter>
      <jdbcPersistenceAdapter dataSource="#oracle-ds"/>
    </persistenceAdapter>
    ...
  </broker>

  <bean id="oracle-ds"
    class="org.apache.commons.dbcp.BasicDataSource"
    destroy-method="close">
    <property name="driverClassName" value="oracle.jdbc.driver.OracleDriver"/>
    <property name="url" value="jdbc:oracle:thin:@localhost:1521:AMQDB"/>
    <property name="username" value="scott"/>
    <property name="password" value="tiger"/>
    <property name="poolPreparedStatements" value="true"/>
  </bean>
</beans>
```

Configuring the clients

Clients of shared JDBC master/slave group must be configured with a failover URL that lists the URLs for all of the brokers in the group.

[Example 3.4 on page 43](#) shows the client failover URL for a group that consists of three brokers: `broker1`, `broker2`, and `broker3`.

Example 3.4. Client URL for a Shared JDBC Master/Slave Group


```
fail
over:(tcp://broker1:61616,tcp://broker2:61616,tcp://broker3:61616)
```

For more information about using the failover protocol see ["Static Failover" on page 15](#).

Reintroducing a failed node

You can restart the failed node at any time and it will rejoin the group. It will rejoin the group as a slave because one of the other brokers already owns the mutex lock on the database table, as shown in [Figure 3.6 on page 43](#).

Figure 3.6. JDBC Master/Slave after Master Restart



Chapter 4. Master/Slave and Broker Networks

Master/slave groups and networks of brokers are very different things. Master/slave groups can be used in a network of brokers to provide fault tolerance to the nodes in the broker network. This requires careful consideration and the use of a special network connection protocol.

Overview

Master/slave groups and broker networks represent different levels of organization. A network of brokers provides a symmetrical group of brokers that share information among all of the members in the group. They are useful for distributing the message processing load among many brokers.

Master/slave groups are asymmetrical> Only one member of the group is active at a time. They are useful for providing fault tolerance when data loss is unacceptable.

You can include a master/slave group as a node in a network of brokers. Using the basic principles of making a master/slave group a node in a broker network, you can scale up to an entire network consisting of master/slave groups.

When combining master/slave groups with broker networks there are two things to remember:

- Network connectors to a master/slave group use a special protocol.
- A broker cannot open a network connection to another member of its master/slave group.

Configuring the connection to a master/slave group

The network connection to a master/slave group needs to do two things:

- Open a connection to the master broker without connecting to the slave brokers.
- Connect to the new master in the case of a failure.

The network connector's reconnect logic will handle the reconnection to the new master in the case of a network failure. The network connector's connection logic, however, attempts to establish connections to all of the specified brokers. To get around the network connector's default behavior,

you use a masterslave URI to specify the list of broker's in the master/slave group. The masterslave URI only allows the connector to connect to one of brokers in the list which will be the master.

The masterslave protocol's URI is a list of the connections points for each broker in the master/slave group. The network connector will traverse the list in order until it establishes a connection.

[Example 4.1 on page 46](#) shows a network connector configured to link to a master/slave group.


Example 4.1. Network Connector to a Master/Slave Group

```
<networkConnectors>
  <networkConnector name="linkToCluster"
    uri="masterslave:(tcp://masterHost:61002,tcp://slaveHost:61002) "
    ... />
</networkConnectors>
```

Host pair with master/slave groups

In order to scale up to a large fault tolerant broker network, it is a good idea to adopt a simple building block as the basis for the network. An effective building block for this purpose is the host pair arrangement shown in [Figure 4.1 on page 46](#).

Figure 4.1. Master/Slave Groups on Two Host Machines




The host pair arrangement consists of two master/slave groups distributed between two host machines. Under normal operating conditions, one master broker is active on each of the two host machines. If one of the machines should fail for some reason, the slave on the other machine takes over, so that you end up with two active brokers on the healthy machine.

When configuring the network connectors, you must remember **not** to open any connectors to brokers in the same group. For example, the network connector for `brokerB1` should be configured to connect to at most `brokerA1` and `brokerA2`.

Network of multiple host pairs

You can easily scale up to a large fault tolerant broker network by adding host pairs, as shown in [Figure 4.2 on page 47](#).

Figure 4.2. Broker Network Consisting of Host Pairs



The preceding network consists of eight master/slave groups distributed over eight host machines. As before, you should open network connectors only to brokers outside the current master/slave group. For example, `brokerA1` can connect to at most the following brokers: `brokerB*`, `brokerC*`, `brokerD*`, `brokerE*`, `brokerF*`, `brokerG*`, and `brokerH*`.

More information

For detailed information on setting up a network of brokers see [Using Networks of Brokers](#).

Index

B

- broker networks
 - master/slave, 45
- broker properties
 - rebalanceClusterClients, 19
 - updateClusterClients, 19
 - updateClusterClientsOnRemove, 19
 - updateClusterFilter, 19

D

- discovery agent
 - Fuse Fabric, 27
 - multicast, 29
 - static, 28
 - zeroconf, 31
- discovery protocol
 - backOffMultiplier, 24
 - initialReconnectDelay, 23
 - maxReconnectAttempts, 24
 - maxReconnectDelay, 23
 - URI, 23
 - useExponentialBackOff, 24
- discovery URI, 23
- discovery://, 23
- discoveryUri, 29, 31
- dynamic failover, 18
 - broker configuration, 19
 - client configuration, 18

F

- fabric://, 27
- failover, 14
 - backOffMultiplier, 16
 - backup, 16
 - broker properties, 19
 - dynamic, 18
 - initialReconnectDelay, 15
 - maxCacheSize, 16
 - maxReconnectAttempts, 16

- maxReconnectDelay, 15
- randomize, 16
- startupMaxReconnectAttempts, 16
- static, 15
- timeout, 16
- trackMessages, 16
- updateURIsSupported, 16
- updateURIsURL, 17
- useExponentialBackOff, 16
- failover URI, 15
 - transport options, 15
- failover://, 15
- Fuse Fabric discovery agent
 - URI, 27

J

- jdbcPersistenceAdapter, 41

M

- master broker
 - reintroduction
 - shared file system, 37
 - shared JDBC, 43
- master/slave
 - broker networks, 45
 - network of brokers, 45
- masterslave, 45
- multicast discovery agent
 - broker configuration, 29
 - URI, 29
- multicast://, 29

N

- network of brokers
 - master/slave, 45
- NFSv3, 34
- NFSv4, 34

O

- OCFS2, 34

P

persistenceAdapter, 36, 41

S

shared file system master/slave

- advantages, 34

- broker configuration, 36, 41

- client configuration, 37

- disadvantages, 34

- incompatible SANs, 34

- initial state, 34

- master failure, 35

- NFSv3, 34

- NFSv4, 34

- OCFS2, 34

- recovery strategies, 35

- reintroducing a node, 37

shared JDBC master/slave

- advantages, 39

- client configuration, 42

- disadvantages, 39

- initial state, 39

- master failure, 40

- recovery strategies, 40

- reintroducing a node, 43

static discovery agent

- URI, 28

static failover, 15

static://, 28

T

transportConnector

- discoveryUri, 29, 31

Z

zeroconf discovery agent

- broker configuration, 31

- URI, 31

zeroconf://, 31