

Cluster configuration

- [Hazelcast configuration](#)
- [Cluster service options](#)
- [Replication configuration](#)
 - [Replication Service configuration](#)
 - [Aeron Media Driver configuration](#)
- [Cluster troubleshooting](#)
 - [Multicast is not enabled on the machine](#)
 - [Using the wrong network interface](#)
 - [When multicast is not available](#)


Hazelcast configuration


The current implementation uses [Hazelcast](#) for resolving cluster nodes. You can find detailed the Hazelcast configuration description on the Hazelcast site, here: [Hazelcast Configuration](#).

If you want to override the default configuration you can:

- provide the `cluster.xml` file with the Hazelcast configuration on your classpath;
- build `com.hazelcast.config.Config` manually and pass it into `HazelcastClusterManager`:

```
Config hazelcastConfig = new Config();
// Now set some stuff on the config (omitted)
ClusterManager mgr = new HazelcastClusterManager(hazelcastConfig);
```

 You can specify a name for the current node with the `instanceName` option. Otherwise, a unique name will be assigned automatically.

 If the quorum size is not specified in the configuration, then the first launched node will be selected as a leader (primary) node.


Cluster service options

In addition to Hazelcast settings, in `HazelcastClusterManager` you can set the following property:

- `timeoutLeaderShutdown` - timeout for waiting for the leader shutdown signal (see [\[Appointment of new leader\]](#)). The default timeout is 90 seconds.

Also, additional properties can be placed in the Advanced Configuration Properties. You can find a detailed description in the [Hazelcast documentation](#).

- `com.epam.fej.cluster.reelectOnLeaderFailure` - this property allows running the leader election process if the previous leader has gone. By default, this property is not specified, which means it has the default value `true` and a new leader will be elected.

 Before the leader has been re-elected, the cluster must contain a member count of greater than or equal to the amount specified by the quorum size property.

Replication configuration

For configuring the replication service, two configuration files are used:

- [Replication Service configuration](#) (*replication.properties*)
- [Aeron Media Driver configuration](#) (*aeron.properties*)

Replication Service configuration

The Replication Service configuration is defined in the `replication.properties` file.

- **fej.replication.leader.sync**
Default (initial) replication mode (synchronous or asynchronous)
Default: false

- **fej.replication.leader.async.timeout**
Default (initial) timeout for synchronous replication in milliseconds. Process can be blocked for this timeout until it receives acknowledgment from the other side.
Default: 0 milliseconds (async mode)
- **fej.replication.leader.receive.buffer.size**
The size of the leader incoming ring buffer. Must be to the power of 2.
Default: 512 bytes
- **fej.replication.leader.receive.wait.strategy**
The wait strategy to use for the leader incoming ring buffer (see [Disruptor User Guide](#)).
Default: com.lmax.disruptor.BlockingWaitStrategy
- **fej.replication.leader.send.buffer.size**
The size of the leader outgoing ring buffer. Must be to the power of 2.
Default: 2048 bytes
- **fej.replication.leader.send.wait.strategy**
The wait strategy used for the leader outgoing ring buffer (see [Disruptor User Guide](#)).
Default: com.lmax.disruptor.BlockingWaitStrategy
- **fej.replication.backup.receive.buffer.size**
The size of the backup incoming ring buffer. Must be to the power of 2.
Default: 1024 bytes
- **fej.replication.backup.receive.wait.strategy**
The wait strategy used for the backup incoming ring buffer (see [Disruptor User Guide](#)).
Default: com.lmax.disruptor.BlockingWaitStrategy
- **fej.replication.backup.send.buffer.size**
The size of the backup outgoing ring buffer. Must be to the power of 2.
Default: 512
- **fej.replication.backup.send.wait.strategy**
The wait strategy used for the backup outgoing ring buffer (see [Disruptor User Guide](#)).
Default: com.lmax.disruptor.BlockingWaitStrategy
- **fej.replication.aeron.mediadriver.embedded**
Use embedded aeron media driver (see [Aeron Embedded Media Driver](#)).
Default: true
- **fej.replication.aeron.idle.strategy**
Provides an IdleStrategy for the thread responsible for communicating with the Aeron Media Driver (see [Aeron Idle Strategies](#)).
Default: uk.co.real_logic.agrona.concurrent.BackoffIdleStrategy

Aeron Media Driver configuration

Please find the description of the Aeron configuration options at its [official page](#).

Cluster troubleshooting

If the default multicast configuration is not working, see below to read about some common causes.

Multicast is not enabled on the machine

It is quite common, in particular on OSX machines, for multicast to be disabled by default. Please google for the answer on how to enable that.

Using the wrong network interface

If you have more than one network interface on your machine (and this can also be the case if you are running VPN software on your machine), then Hazelcast may be using the wrong one.

To tell Hazelcast to use a specific interface, you can provide the IP address of the interface in the `interfaces` element of the configuration. Make sure you set the `enabled` attribute to `true`. For example:

```
<interfaces enabled="true">
  <interface>192.168.1.20</interface>
</interfaces>
```

When multicast is not available

In some cases, you may not be able to use multicast as it might not be available in your environment. In that case, you should configure another transport, for example, TCP to use TCP sockets, or AWS when running on Amazon EC2.

For more information on available Hazelcast transports and how to configure them, please consult the [Hazelcast Configuration](#).