



# **Red Hat JBoss Data Grid 6.5**

## **Administration and Configuration Guide**

For use with Red Hat JBoss Data Grid 6.5.1



# Red Hat JBoss Data Grid 6.5 Administration and Configuration Guide

---

For use with Red Hat JBoss Data Grid 6.5.1

Misha Husnain Ali  
Red Hat Engineering Content Services  
mhusnain@redhat.com

Gemma Sheldon  
Red Hat Engineering Content Services  
gsheldon@redhat.com

Rakesh Ghatvisave  
Red Hat Engineering Content Services  
rghatvis@redhat.com

Christian Huffman  
Red Hat Engineering Content Services  
chuffman@redhat.com

## Legal Notice

Copyright © 2016 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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 guide presents information about the administration and configuration of Red Hat JBoss Data Grid 6.5.1

## Table of Contents

<b>CHAPTER 1. SETTING UP RED HAT JBOSS DATA GRID</b> .....	<b>8</b>
1.1. PREREQUISITES	8
1.2. STEPS TO SET UP RED HAT JBOSS DATA GRID	8
<b>PART I. SET UP A CACHE MANAGER</b> .....	<b>11</b>
<b>CHAPTER 2. CACHE MANAGERS</b> .....	<b>12</b>
2.1. TYPES OF CACHE MANAGERS	12
2.2. CREATING CACHEMANAGERS	12
2.3. MULTIPLE CACHE MANAGERS	13
<b>PART II. SET UP JVM MEMORY MANAGEMENT</b> .....	<b>15</b>
<b>CHAPTER 3. SET UP EVICTION</b> .....	<b>16</b>
3.1. ABOUT EVICTION	16
3.2. EVICTION STRATEGIES	16
3.3. USING EVICTION	17
<b>CHAPTER 4. SET UP EXPIRATION</b> .....	<b>20</b>
4.1. ABOUT EXPIRATION	20
4.2. EXPIRATION OPERATIONS	20
4.3. EVICTION AND EXPIRATION COMPARISON	20
4.4. CACHE ENTRY EXPIRATION NOTIFICATIONS	21
4.5. CONFIGURE EXPIRATION	21
4.6. MORTAL AND IMMORTAL DATA	21
4.7. TROUBLESHOOTING EXPIRATION	22
<b>PART III. MONITOR YOUR CACHE</b> .....	<b>23</b>
<b>CHAPTER 5. SET UP LOGGING</b> .....	<b>24</b>
5.1. ABOUT LOGGING	24
5.2. SUPPORTED APPLICATION LOGGING FRAMEWORKS	24
5.3. BOOT LOGGING	25
5.4. LOGGING ATTRIBUTES	25
5.5. LOGGING SAMPLE CONFIGURATIONS	31
<b>PART IV. SET UP CACHE MODES</b> .....	<b>37</b>
<b>CHAPTER 6. CACHE MODES</b> .....	<b>38</b>
6.1. ABOUT CACHE CONTAINERS	38
6.2. LOCAL MODE	39
6.3. CLUSTERED MODES	40
6.4. STATE TRANSFER	42
<b>CHAPTER 7. SET UP DISTRIBUTION MODE</b> .....	<b>44</b>
7.1. ABOUT DISTRIBUTION MODE	44
7.2. DISTRIBUTION MODE'S CONSISTENT HASH ALGORITHM	44
7.3. LOCATING ENTRIES IN DISTRIBUTION MODE	44
7.4. RETURN VALUES IN DISTRIBUTION MODE	45
7.5. CONFIGURE DISTRIBUTION MODE (REMOTE CLIENT-SERVER MODE)	45
7.6. CONFIGURE DISTRIBUTION MODE (LIBRARY MODE)	46
7.7. SYNCHRONOUS AND ASYNCHRONOUS DISTRIBUTION	47
7.8. GET AND PUT USAGE IN DISTRIBUTION MODE	47

<b>CHAPTER 8. SET UP REPLICATION MODE</b> .....	<b>49</b>
8.1. ABOUT REPLICATION MODE	49
8.2. OPTIMIZED REPLICATION MODE USAGE	49
8.3. CONFIGURE REPLICATION MODE (REMOTE CLIENT-SERVER MODE)	49
8.4. CONFIGURE REPLICATION MODE (LIBRARY MODE)	50
8.5. SYNCHRONOUS AND ASYNCHRONOUS REPLICATION	51
8.6. THE REPLICATION QUEUE	52
8.7. ABOUT REPLICATION GUARANTEES	53
8.8. REPLICATION TRAFFIC ON INTERNAL NETWORKS	53
<b>CHAPTER 9. SET UP INVALIDATION MODE</b> .....	<b>55</b>
9.1. ABOUT INVALIDATION MODE	55
9.2. CONFIGURE INVALIDATION MODE (REMOTE CLIENT-SERVER MODE)	55
9.3. CONFIGURE INVALIDATION MODE (LIBRARY MODE)	56
9.4. SYNCHRONOUS/ASYNCHRONOUS INVALIDATION	57
9.5. THE L1 CACHE AND INVALIDATION	57
<b>PART V. REMOTE CLIENT-SERVER MODE INTERFACES</b> .....	<b>58</b>
<b>CHAPTER 10. THE ASYNCHRONOUS API</b> .....	<b>59</b>
10.1. ASYNCHRONOUS API BENEFITS	59
10.2. ABOUT ASYNCHRONOUS PROCESSES	59
10.3. RETURN VALUES AND THE ASYNCHRONOUS API	60
<b>CHAPTER 11. THE REST INTERFACE</b> .....	<b>61</b>
11.1. RUBY CLIENT CODE	61
11.2. USING JSON WITH RUBY EXAMPLE	61
11.3. PYTHON CLIENT CODE	62
11.4. JAVA CLIENT CODE	62
11.5. THE REST INTERFACE CONNECTOR	65
11.6. USING THE REST INTERFACE	66
11.7. REST INTERFACE SECURITY	71
<b>CHAPTER 12. THE MEMCACHED INTERFACE</b> .....	<b>74</b>
12.1. ABOUT MEMCACHED SERVERS	74
12.2. MEMCACHED STATISTICS	74
12.3. THE MEMCACHED INTERFACE CONNECTOR	76
12.4. MEMCACHED INTERFACE SECURITY	77
<b>CHAPTER 13. THE HOT ROD INTERFACE</b> .....	<b>79</b>
13.1. ABOUT HOT ROD	79
13.2. THE BENEFITS OF USING HOT ROD OVER MEMCACHED	79
13.3. HOT ROD HASH FUNCTIONS	80
13.4. THE HOT ROD INTERFACE CONNECTOR	80
13.5. HOT ROD HEADERS	82
13.6. HOT ROD OPERATIONS	88
13.7. HOT ROD OPERATION VALUES	105
13.8. PUT REQUEST EXAMPLE	109
13.9. HOT ROD JAVA CLIENT	110
13.10. HOT ROD C++ CLIENT	114
13.11. HOT ROD C# CLIENT	117
13.12. INTEROPERABILITY BETWEEN HOT ROD C++ AND HOT ROD JAVA CLIENT	123
13.13. COMPATIBILITY BETWEEN SERVER AND HOT ROD CLIENT VERSIONS	124
<b>PART VI. SET UP LOCKING FOR THE CACHE</b> .....	<b>126</b>

<b>CHAPTER 14. LOCKING</b> .....	<b>127</b>
14.1. CONFIGURE LOCKING (REMOTE CLIENT-SERVER MODE)	127
14.2. CONFIGURE LOCKING (LIBRARY MODE)	127
14.3. LOCKING TYPES	128
14.4. LOCKING OPERATIONS	130
<b>CHAPTER 15. SET UP LOCK STRIPING</b> .....	<b>132</b>
15.1. ABOUT LOCK STRIPING	132
15.2. CONFIGURE LOCK STRIPING (REMOTE CLIENT-SERVER MODE)	132
15.3. CONFIGURE LOCK STRIPING (LIBRARY MODE)	132
<b>CHAPTER 16. SET UP ISOLATION LEVELS</b> .....	<b>134</b>
16.1. ABOUT ISOLATION LEVELS	134
16.2. ABOUT READ_COMMITTED	134
16.3. ABOUT REPEATABLE_READ	134
<b>PART VII. SET UP AND CONFIGURE A CACHE STORE</b> .....	<b>136</b>
<b>CHAPTER 17. CACHE STORES</b> .....	<b>137</b>
17.1. CACHE LOADERS AND CACHE WRITERS	137
17.2. CACHE STORE CONFIGURATION	137
17.3. SHARED CACHE STORES	140
17.4. CONNECTION FACTORIES	141
<b>CHAPTER 18. CACHE STORE IMPLEMENTATIONS</b> .....	<b>142</b>
18.1. CACHE STORE COMPARISON	142
18.2. CACHE STORE CONFIGURATION DETAILS (LIBRARY MODE)	142
18.3. CACHE STORE CONFIGURATION DETAILS (REMOTE CLIENT-SERVER MODE)	147
18.4. SINGLE FILE CACHE STORE	149
18.5. LEVELDB CACHE STORE	151
18.6. JDBC BASED CACHE STORES	155
18.7. THE REMOTE CACHE STORE	166
18.8. JPA CACHE STORE	167
18.9. CUSTOM CACHE STORES	170
<b>PART VIII. SET UP PASSIVATION</b> .....	<b>176</b>
<b>CHAPTER 19. ACTIVATION AND PASSIVATION MODES</b> .....	<b>177</b>
19.1. PASSIVATION MODE BENEFITS	177
19.2. CONFIGURE PASSIVATION	177
19.3. EVICTION AND PASSIVATION	178
<b>PART IX. SET UP CACHE WRITING</b> .....	<b>180</b>
<b>CHAPTER 20. CACHE WRITING MODES</b> .....	<b>181</b>
20.1. WRITE-THROUGH CACHING	181
20.2. WRITE-BEHIND CACHING	182
<b>PART X. MONITOR CACHES AND CACHE MANAGERS</b> .....	<b>185</b>
<b>CHAPTER 21. SET UP JAVA MANAGEMENT EXTENSIONS (JMX)</b> .....	<b>186</b>
21.1. ABOUT JAVA MANAGEMENT EXTENSIONS (JMX)	186
21.2. USING JMX WITH RED HAT JBOSS DATA GRID	186
21.3. JMX STATISTIC LEVELS	186
21.4. ENABLE JMX FOR CACHE INSTANCES	186
21.5. ENABLE JMX FOR CACHEMANAGERS	187

---

21.6. DISABLING THE CACHESTORE VIA JMX WHEN USING ROLLING UPGRADES	187
21.7. MULTIPLE JMX DOMAINS	187
21.8. MBEANS	188
<b>CHAPTER 22. SET UP JBOSS OPERATIONS NETWORK (JON)</b>	<b>190</b>
22.1. ABOUT JBOSS OPERATIONS NETWORK (JON)	190
22.2. DOWNLOAD JBOSS OPERATIONS NETWORK (JON)	190
22.3. JBOSS OPERATIONS NETWORK SERVER INSTALLATION	192
22.4. JBOSS OPERATIONS NETWORK AGENT	192
22.5. JBOSS OPERATIONS NETWORK FOR REMOTE CLIENT-SERVER MODE	192
22.6. JBOSS OPERATIONS NETWORK REMOTE-CLIENT SERVER PLUGIN	193
22.7. JBOSS OPERATIONS NETWORK FOR LIBRARY MODE	199
22.8. JBOSS OPERATIONS NETWORK PLUG-IN QUICKSTART	211
22.9. OTHER MANAGEMENT TOOLS AND OPERATIONS	211
<b>PART XI. COMMAND LINE TOOLS</b>	<b>213</b>
<b>CHAPTER 23. RED HAT JBOSS DATA GRID CLIS</b>	<b>214</b>
23.1. RED HAT JBOSS DATA GRID LIBRARY MODE CLI	214
23.2. RED HAT DATA GRID SERVER CLI	215
23.3. CLI COMMANDS	216
<b>PART XII. OTHER RED HAT JBOSS DATA GRID FUNCTIONS</b>	<b>224</b>
<b>CHAPTER 24. SET UP THE L1 CACHE</b>	<b>225</b>
24.1. ABOUT THE L1 CACHE	225
24.2. L1 CACHE CONFIGURATION	225
<b>CHAPTER 25. SET UP TRANSACTIONS</b>	<b>227</b>
25.1. ABOUT TRANSACTIONS	227
25.2. CONFIGURE TRANSACTIONS	228
25.3. TRANSACTION RECOVERY	231
25.4. DEADLOCK DETECTION	232
<b>CHAPTER 26. CONFIGURE JGROUPS</b>	<b>234</b>
26.1. CONFIGURE RED HAT JBOSS DATA GRID INTERFACE BINDING (REMOTE CLIENT-SERVER MODE)	234
26.2. CONFIGURE JGROUPS (LIBRARY MODE)	237
26.3. TEST MULTICAST USING JGROUPS	241
<b>CHAPTER 27. USE RED HAT DATA GRID WITH AMAZON WEB SERVICES</b>	<b>246</b>
27.1. THE S3_PING JGROUPS DISCOVERY PROTOCOL	246
27.2. S3_PING CONFIGURATION OPTIONS	246
<b>CHAPTER 28. USE RED HAT JBOSS DATA GRID WITH GOOGLE COMPUTE ENGINE</b>	<b>251</b>
28.1. THE GOOGLE_PING PROTOCOL	251
28.2. GOOGLE_PING CONFIGURATION	251
<b>CHAPTER 29. HIGH AVAILABILITY USING SERVER HINTING</b>	<b>253</b>
29.1. ESTABLISHING SERVER HINTING WITH JGROUPS	253
29.2. CONFIGURE SERVER HINTING (REMOTE CLIENT-SERVER MODE)	253
29.3. CONFIGURE SERVER HINTING (LIBRARY MODE)	254
29.4. CONSISTENTHASHFACTORIES	254
29.5. KEY AFFINITY SERVICE	256
<b>CHAPTER 30. SET UP CROSS-DATACENTER REPLICATION</b>	<b>258</b>



30.1. CROSS-DATACENTER REPLICATION OPERATIONS	258
30.2. CONFIGURE CROSS-DATACENTER REPLICATION	259
30.3. TAKING A SITE OFFLINE	265
30.4. STATE TRANSFER BETWEEN SITES	267
30.5. CONFIGURE MULTIPLE SITE MASTERS	270
<b>CHAPTER 31. EXTERNALIZE SESSIONS</b>	<b>273</b>
31.1. EXTERNALIZE HTTP SESSION FROM JBOSS EAP 6.4 AND LATER TO JBOSS DATA GRID	273
<b>CHAPTER 32. HANDLING NETWORK PARTITIONS (SPLIT BRAIN)</b>	<b>275</b>
32.1. DETECTING AND RECOVERING FROM A SPLIT-BRAIN PROBLEM	275
32.2. SPLIT BRAIN TIMING: DETECTING A SPLIT	278
32.3. SPLIT BRAIN TIMING: RECOVERING FROM A SPLIT	278
32.4. DETECTING AND RECOVERING FROM SUCCESSIVE CRASHED NODES	279
32.5. NETWORK PARTITION RECOVERY EXAMPLES	279
32.6. CONFIGURE PARTITION HANDLING	294
<b>APPENDIX A. RECOMMENDED JGROUPS VALUES FOR JBOSS DATA GRID</b>	<b>296</b>
A.1. SUPPORTED JGROUPS PROTOCOLS	296
A.2. TCP DEFAULT AND RECOMMENDED VALUES	300
A.3. UDP DEFAULT AND RECOMMENDED VALUES	307
A.4. JBOSS DATA GRID JGROUPS CONFIGURATION FILES	314
<b>APPENDIX B. CONNECTING WITH JCONSOLE</b>	<b>315</b>
B.1. CONNECT TO JDG VIA JCONSOLE	315
<b>APPENDIX C. JMX MBEANS IN REDHAT JBOSS DATA GRID</b>	<b>319</b>
C.1. ACTIVATION	319
C.2. CACHE	319
C.3. CACHECONTAINERSTATS	320
C.4. CACHELOADER	321
C.5. CACHEMANAGER	322
C.6. CACHESTORE	324
C.7. CLUSTERCACHESTATS	324
C.8. DEADLOCKDETECTINGLOCKMANAGER	327
C.9. DISTRIBUTIONMANAGER	328
C.10. INTERPRETER	328
C.11. INVALIDATION	329
C.12. LOCKMANAGER	329
C.13. LOCALTOPOLOGYMANAGER	330
C.14. MASSINDEXER	331
C.15. PASSIVATION	331
C.16. RECOVERYADMIN	332
C.17. ROLLINGUPGRADEMANAGER	332
C.18. RPCMANAGER	333
C.19. STATETRANFERMANAGER	334
C.20. STATISTICS	335
C.21. TRANSACTIONS	336
C.22. TRANSPORT	337
C.23. XSITEADMIN	338
<b>APPENDIX D. CONFIGURATION RECOMMENDATIONS</b>	<b>341</b>
D.1. TIMEOUT VALUES	341
<b>APPENDIX E. PERFORMANCE RECOMMENDATIONS</b>	<b>342</b>

E.1. CONCURRENT STARTUP FOR LARGE CLUSTERS	342
<b>APPENDIX F. REFERENCES</b> .....	<b>343</b>
F.1. ABOUT CONSISTENCY	343
F.2. ABOUT CONSISTENCY GUARANTEE	343
F.3. ABOUT JBOSS CACHE	343
F.4. ABOUT RELAY2	343
F.5. ABOUT RETURN VALUES	344
F.6. ABOUT RUNNABLE INTERFACES	344
F.7. ABOUT TWO PHASE COMMIT (2PC)	344
F.8. ABOUT KEY-VALUE PAIRS	344
F.9. THE EXTERNALIZER	344
F.10. HASH SPACE ALLOCATION	345
<b>APPENDIX G. REVISION HISTORY</b> .....	<b>348</b>



# CHAPTER 1. SETTING UP RED HAT JBOSS DATA GRID

## 1.1. PREREQUISITES

The only prerequisites to set up Red Hat JBoss Data Grid is a Java Virtual Machine and that the most recent supported version of the product is installed on your system.

[Report a bug](#)

## 1.2. STEPS TO SET UP RED HAT JBOSS DATA GRID

The following steps outline the necessary (and optional, where stated) steps for a first time basic configuration of Red Hat JBoss Data Grid. It is recommended that the steps are followed in the order specified and not skipped unless they are identified as optional steps.

### Procedure 1.1. Set Up JBoss Data Grid

#### 1. Set Up the Cache Manager

The first step in a JBoss Data Grid configuration is a cache manager. Cache managers can retrieve cache instances and create cache instances quickly and easily using previously specified configuration templates. For details about setting up a cache manager, see [Part I, “Set Up a Cache Manager”](#)

#### 2. Set Up JVM Memory Management

An important step in configuring your JBoss Data Grid is to set up memory management for your Java Virtual Machine (JVM). JBoss Data Grid offers features such as eviction and expiration to help manage the JVM memory.

##### a. Set Up Eviction

Use eviction to specify the logic used to remove entries from the in-memory cache implementation based on how often they are used. JBoss Data Grid offers different eviction strategies for finer control over entry eviction in your data grid. Eviction strategies and instructions to configure them are available in [Chapter 3, Set Up Eviction](#).

##### b. Set Up Expiration

To set upper limits to an entry's time in the cache, attach expiration information to each entry. Use expiration to set up the maximum period an entry is allowed to remain in the cache and how long the retrieved entry can remain idle before being removed from the cache. For details, see [Chapter 4, Set Up Expiration](#)

#### 3. Monitor Your Cache

JBoss Data Grid uses logging via JBoss Logging to help users monitor their caches.

##### a. Set Up Logging

It is not mandatory to set up logging for your JBoss Data Grid, but it is highly recommended. JBoss Data Grid uses JBoss Logging, which allows the user to easily set up automated logging for operations in the data grid. Logs can subsequently be used to troubleshoot errors and identify the cause of an unexpected failure. For details, see [Chapter 5, Set Up Logging](#)

#### 4. Set Up Cache Modes

Cache modes are used to specify whether a cache is local (simple, in-memory cache) or a clustered cache (replicates state changes over a small subset of nodes). Additionally, if a cache is clustered, either replication, distribution or invalidation mode must be applied to determine

how the changes propagate across the subset of nodes. For details, see [Part IV, “Set Up Cache Modes”](#)

## 5. Set Up Locking for the Cache

When replication or distribution is in effect, copies of entries are accessible across multiple nodes. As a result, copies of the data can be accessed or modified concurrently by different threads. To maintain consistency for all copies across nodes, configure locking. For details, see [Part VI, “Set Up Locking for the Cache”](#) and [Chapter 16, \*Set Up Isolation Levels\*](#)

## 6. Set Up and Configure a Cache Store

JBoss Data Grid offers the passivation feature (or cache writing strategies if passivation is turned off) to temporarily store entries removed from memory in a persistent, external cache store. To set up passivation or a cache writing strategy, you must first set up a cache store.

### a. Set Up a Cache Store

The cache store serves as a connection to the persistent store. Cache stores are primarily used to fetch entries from the persistent store and to push changes back to the persistent store. For details, see [Part VII, “Set Up and Configure a Cache Store”](#)

### b. Set Up Passivation

Passivation stores entries evicted from memory in a cache store. This feature allows entries to remain available despite not being present in memory and prevents potentially expensive write operations to the persistent cache. For details, see [Part VIII, “Set Up Passivation”](#)

### c. Set Up a Cache Writing Strategy

If passivation is disabled, every attempt to write to the cache results in writing to the cache store. This is the default Write-Through cache writing strategy. Set the cache writing strategy to determine whether these cache store writes occur synchronously or asynchronously. For details, see [Part IX, “Set Up Cache Writing”](#)

## 7. Monitor Caches and Cache Managers

JBoss Data Grid includes two primary tools to monitor the cache and cache managers once the data grid is up and running.

### a. Set Up JMX

JMX is the standard statistics and management tool used for JBoss Data Grid. Depending on the use case, JMX can be configured at a cache level or a cache manager level or both. For details, see [Chapter 21, \*Set Up Java Management Extensions \(JMX\)\*](#)

### b. Set Up Red Hat JBoss Operations Network (JON)

Red Hat JBoss Operations Network (JON) is the second monitoring solution available for JBoss Data Grid. JBoss Operations Network (JON) offers a graphical interface to monitor runtime parameters and statistics for caches and cache managers. For details, see [Chapter 22, \*Set Up JBoss Operations Network \(JON\)\*](#)

## 8. Introduce Topology Information

Optionally, introduce topology information to your data grid to specify where specific types of information or objects in your data grid are located. Server hinting is one of the ways to introduce topology information in JBoss Data Grid.

### a. Set Up Server Hinting

When set up, server hinting provides high availability by ensuring that the original and backup copies of data are not stored on the same physical server, rack or data center. This is optional in cases such as a replicated cache, where all data is backed up on all servers, racks and data centers. For details, see [Chapter 29, \*High Availability Using Server Hinting\*](#)

The subsequent chapters detail each of these steps towards setting up a standard JBoss Data Grid configuration.

[Report a bug](#)

## PART I. SET UP A CACHE MANAGER

## CHAPTER 2. CACHE MANAGERS

A Cache Manager is the primary mechanism to retrieve a cache instance in Red Hat JBoss Data Grid, and is a starting point for using the cache.

In JBoss Data Grid, a cache manager is useful because:

- it creates multiple cache instances on demand using a provided standard.
- it retrieves existing cache instances (i.e. caches that have already been created).

[Report a bug](#)

### 2.1. TYPES OF CACHE MANAGERS

Red Hat JBoss Data Grid offers the following Cache Managers:

- **EmbeddedCacheManager** is a cache manager that runs within the Java Virtual Machine (JVM) used by the client. Currently, JBoss Data Grid offers only the **DefaultCacheManager** implementation of the **EmbeddedCacheManager** interface.
- **RemoteCacheManager** is used to access remote caches. When started, the **RemoteCacheManager** instantiates connections to the Hot Rod server (or multiple Hot Rod servers). It then manages the persistent **TCP** connections while it runs. As a result, **RemoteCacheManager** is resource-intensive. The recommended approach is to have a single **RemoteCacheManager** instance for each Java Virtual Machine (JVM).

[Report a bug](#)

### 2.2. CREATING CACHEMANAGERS

#### 2.2.1. Create a New RemoteCacheManager

##### Procedure 2.1. Configure a New RemoteCacheManager

```
import org.infinispan.client.hotrod.configuration.Configuration;
import org.infinispan.client.hotrod.configuration.ConfigurationBuilder;

Configuration conf = new
ConfigurationBuilder().addServer().host("localhost").port(11222).build();
RemoteCacheManager manager = new RemoteCacheManager(conf);
RemoteCache defaultCache = manager.getCache();
```

1. Use the **ConfigurationBuilder()** constructor to create a new configuration builder. The **.addServer()** method adds a remote server, configured via the **.host(<hostname|ip>)** and **.port(<port>)** properties.
2. Create a new **RemoteCacheManager** using the supplied configuration.
3. Retrieve the default cache from the remote server.

[Report a bug](#)



## 2.2.2. Create a New Embedded Cache Manager

Use the following instructions to create a new `EmbeddedCacheManager` without using CDI:

### Procedure 2.2. Create a New Embedded Cache Manager

1. Create a configuration XML file. For example, create the `my-config-file.xml` file on the classpath (in the `resources/` folder) and add the configuration information in this file.
2. Use the following programmatic configuration to create a cache manager using the configuration file:

```
EmbeddedCacheManager manager = new DefaultCacheManager("my-config-
file.xml");
Cache defaultCache = manager.getCache();
```

The outlined procedure creates a new `EmbeddedCacheManager` using the configuration specified in `my-config-file.xml`.

[Report a bug](#)

## 2.2.3. Create a New Embedded Cache Manager Using CDI

Use the following steps to create a new `EmbeddedCacheManager` instance using CDI:

### Procedure 2.3. Use CDI to Create a New EmbeddedCacheManager

1. Specify a default configuration:

```
public class Config
    @Produces
    public EmbeddedCacheManager defaultCacheManager() {
        ConfigurationBuilder builder = new ConfigurationBuilder();
        Configuration configuration =
builder.eviction().strategy(EvictionStrategy.LRU).maxEntries(100).bu
ild();
        return new DefaultCacheManager(configuration);
    }
}
```

2. Inject the default cache manager.

```
<!-- Additional configuration information here -->
@Inject
EmbeddedCacheManager cacheManager;
<!-- Additional configuration information here -->
```

[Report a bug](#)

## 2.3. MULTIPLE CACHE MANAGERS

Cache managers are an entry point to the cache and Red Hat JBoss Data Grid allows users to create multiple cache managers. Each cache manager is configured with a different global configuration, which includes settings for things like JMX, executors and clustering.

[Report a bug](#)

### 2.3.1. Create Multiple Caches with a Single Cache Manager

Red Hat JBoss Data Grid allows using the same cache manager to create multiple caches, each with a different cache mode (synchronous and asynchronous cache modes).

[Report a bug](#)

### 2.3.2. Using Multiple Cache Managers

Red Hat JBoss Data Grid allows multiple cache managers to be used. In most cases, such as with replication and networking components, cache instances share internal components and a single cache manager is sufficient.

However, if multiple caches are required to have different network characteristics, for example if one cache uses the **TCP** protocol and the other uses the **UDP** protocol, multiple cache managers must be used.

[Report a bug](#)

### 2.3.3. Create Multiple Cache Managers

Red Hat JBoss Data Grid allows users to create multiple cache managers of various types by repeating the procedure used to create the first cache manager (and adjusting the contents of the configuration file, if required).

To use the declarative API to create multiple new cache managers, copy the contents of the **infinispan.xml** file to a new configuration file. Edit the new file for the desired configuration and then use the new file for a new cache manager.

[Report a bug](#)

## PART II. SET UP JVM MEMORY MANAGEMENT

## CHAPTER 3. SET UP EVICTION

### 3.1. ABOUT EVICTION

Eviction is the process of removing entries from memory to prevent running out of memory. Entries that are evicted from memory remain in cache stores and the rest of the cluster to prevent permanent data loss.

Red Hat JBoss Data Grid executes eviction tasks by utilizing user threads which are already interacting with the data container. JBoss Data Grid uses a separate thread to prune expired cache entries from the cache.

Eviction occurs individually on a per node basis, rather than occurring as a cluster-wide operation. Each node uses an eviction thread to analyze the contents of its in-memory container to determine which entries require eviction. The free memory in the Java Virtual Machine (JVM) is not a consideration during the eviction analysis, even as a threshold to initialize entry eviction.

In JBoss Data Grid, eviction provides a mechanism to efficiently remove entries from the in-memory representation of a cache and push them to a persistent store. This ensures that the memory can always accommodate new entries as they are fetched and that evicted entries are preserved in the cluster instead of lost.

Additionally, eviction strategies can be used as required for your configuration to set up which entries are evicted and when eviction occurs.

#### See Also:

- [Section 4.3, “Eviction and Expiration Comparison”](#)

[Report a bug](#)

### 3.2. EVICTION STRATEGIES

Each eviction strategy has specific benefits and use cases, as outlined below:

**Table 3.1. Eviction Strategies**

Strategy Name	Operations	Details
<code>EvictionStrategy.NONE</code>	No eviction occurs.	This is the default eviction strategy in Red Hat JBoss Data Grid.
<code>EvictionStrategy.LRU</code>	Least Recently Used eviction strategy. This strategy evicts entries that have not been used for the longest period. This ensures that entries that are reused periodically remain in memory.	

Strategy Name	Operations	Details
<b>EvictionStrategy.UNORDERED</b>	Unordered eviction strategy. This strategy evicts entries without any ordered algorithm and may therefore evict entries that are required later. However, this strategy saves resources because no algorithm related calculations are required before eviction.	This strategy is recommended for testing purposes and not for a real work implementation.
<b>EvictionStrategy.LIRS</b>	Low Inter-Reference Recency Set eviction strategy.	LIRS is an eviction algorithm that suits a large variety of production use cases.

[Report a bug](#)

### 3.2.1. LRU Eviction Algorithm Limitations

In the Least Recently Used (LRU) eviction algorithm, the least recently used entry is evicted first. The entry that has not been accessed the longest gets evicted first from the cache. However, LRU eviction algorithm sometimes does not perform optimally in cases of weak access locality. The weak access locality is a technical term used for entries which are put in the cache and not accessed for a long time and entries to be accessed soonest are replaced. In such cases, problems such as the following can appear:

- Single use access entries are not replaced in time.
- Entries that are accessed first are unnecessarily replaced.

[Report a bug](#)

## 3.3. USING EVICTION

In Red Hat JBoss Data Grid, eviction is disabled by default. If an empty `<eviction />` element is used to enable eviction without any strategy or maximum entries settings, the following default values are automatically implemented:

- Strategy: If no eviction strategy is specified, **EvictionStrategy.NONE** is assumed as a default.
- max-entries/maxEntries: If no value is specified, the **max-entries**/maxEntries value is set to **-1**, which allows unlimited entries.

[Report a bug](#)

### 3.3.1. Initialize Eviction

To initialize eviction, set the eviction element's **max-entries** attributes value to a number greater than zero. Adjust the value set for **max-entries** to discover the optimal value for your configuration. It is important to remember that if too large a value is set for **max-entries**, Red Hat JBoss Data Grid runs

out of memory.

The following procedure outlines the steps to initialize eviction in JBoss Data Grid:

### Procedure 3.1. Initialize Eviction

#### 1. Add the Eviction Tag

Add the `<eviction>` tag to your project's `<cache>` tags as follows:

```
<eviction />
```

#### 2. Set the Eviction Strategy

Set the *strategy* value to set the eviction strategy employed. Possible values are **LRU**, **UNORDERED** and **LIRS** (or **NONE** if no eviction is required). The following is an example of this step:

```
<eviction strategy="LRU" />
```

#### 3. Set the Maximum Entries

Set the maximum number of entries allowed in memory. The default value is **-1** for unlimited entries.

a. In Library mode, set the *maxEntries* parameter as follows:

```
<eviction strategy="LRU" maxEntries="200" />
```

b. In Remote Client Server mode, set the *max-entries* as follows:

```
<eviction strategy="LRU" max-entries="200" />
```

### Result

Eviction is configured for the target cache.

[Report a bug](#)

### 3.3.2. Eviction Configuration Examples

Configure eviction in Red Hat JBoss Data Grid using the configuration bean or the XML file. Eviction configuration is done on a per-cache basis.

- A sample XML configuration for Library mode is as follows:

```
<eviction strategy="LRU" maxEntries="2000"/>
```

- A sample XML configuration for Remote Client Server Mode is as follows:

```
<eviction strategy="LRU" max-entries="20"/>
```

- A sample programmatic configuration for Library Mode is as follows:

```
Configuration c = new
```

```
ConfigurationBuilder().eviction().strategy(EvictionStrategy.LRU)
    .maxEntries(2000)
    .build();
```



#### NOTE

JBoss Data Grid's Library mode uses the *maxEntries* parameter while Remote Client-Server mode uses the *max-entries* parameter to configure eviction.

[Report a bug](#)

### 3.3.3. Eviction Configuration Troubleshooting

In Red Hat JBoss Data Grid, the size of a cache can be larger than the value specified for the *max-entries* parameter of the *eviction* element. This is because although the *max-entries* value can be configured to a value that is not a power of two, the underlying algorithm will alter the value to **V**, where **V** is the closest power of two value that is larger than the *max-entries* value. Eviction algorithms are in place to ensure that the size of the cache container will never exceed the value **V**.

[Report a bug](#)

### 3.3.4. Eviction and Passivation

To ensure that a single copy of an entry remains, either in memory or in a cache store, use passivation in conjunction with eviction.

The primary reason to use passivation instead of a normal cache store is that updating entries require less resources when passivation is in use. This is because passivation does not require an update to the cache store.

[Report a bug](#)

## CHAPTER 4. SET UP EXPIRATION

### 4.1. ABOUT EXPIRATION

Red Hat JBoss Data Grid uses expiration to attach one or both of the following values to an entry:

- A lifespan value.
- A maximum idle time value.

Expiration can be specified on a per-entry or per-cache basis and the per-entry configuration overrides per-cache configurations. If expiration is not configured at the cache level, cache entries are created immortal (i.e. they will never expire) as a default. Conversely, if expiration is configured at the cache level, the expiration defaults apply to all entries which do not explicitly specify a ***lifespan*** or ***maxIdle*** value.

Expired entries, unlike evicted entries, are removed globally, which removes them from memory, cache stores and the cluster.

Expiration automates the removal of entries that have not been used for a specified period of time from the memory. Expiration and eviction are different because:

- expiration removes entries based on the period they have been in memory. Expiration only removes entries when the life span period concludes or when an entry has been idle longer than the specified idle time.
- eviction removes entries based on how recently (and often) they are used. Eviction only removes entries when too many entries are present in the memory. If a cache store has been configured, evicted entries are persisted in the cache store.

[Report a bug](#)

### 4.2. EXPIRATION OPERATIONS

Expiration in Red Hat JBoss Data Grid allows you to set a life span or maximum idle time value for each key/value pair stored in the cache.

The life span or maximum idle time can be set to apply cache-wide or defined for each key/value pair using the cache API. The life span (***lifespan***) or maximum idle time (***maxIdle*** in Library Mode and ***max-idle*** in Remote Client-Server Mode) defined for an individual key/value pair overrides the cache-wide default for the entry in question.

[Report a bug](#)

### 4.3. EVICTION AND EXPIRATION COMPARISON

Expiration is a top-level construct in Red Hat JBoss Data Grid, and is represented in the global configuration, as well as the cache API.

Eviction is limited to the cache instance it is used in, whilst expiration is cluster-wide. Expiration life spans (***lifespan***) and idle time (***maxIdle*** in Library Mode and ***max-idle*** in Remote Client-Server Mode) values are replicated alongside each cache entry.

[Report a bug](#)



## 4.4. CACHE ENTRY EXPIRATION NOTIFICATIONS

Red Hat JBoss Data Grid does not guarantee that an eviction occurs immediately upon timeout. Instead, a number of mechanisms are used in collaboration to ensure efficient eviction. An expired entry is removed from the cache when either:

- An entry is passivated/overflowed to disk and is discovered to have expired.
- The eviction maintenance thread discovers that an entry it has found is expired.

If a user requests an entry that is expired but not yet evicted, a null value is sent to the user. This mechanism ensures that the user never receives an expired entry. The entry is eventually removed by the eviction thread.

[Report a bug](#)

## 4.5. CONFIGURE EXPIRATION

In Red Hat JBoss Data Grid, expiration is configured in a manner similar to eviction.

### Procedure 4.1. Configure Expiration

#### 1. Add the Expiration Tag

Add the `<expiration>` tag to your project's `<cache>` tags as follows:

```
<expiration />
```

#### 2. Set the Expiration Lifespan

Set the ***lifespan*** value to set the period of time (in milliseconds) an entry can remain in memory. The following is an example of this step:

```
<expiration lifespan="1000" />
```

#### 3. Set the Maximum Idle Time

Set the time that entries are allowed to remain idle (unused) after which they are removed (in milliseconds). The default value is `-1` for unlimited time.

- In Library mode, set the ***maxIdle*** parameter as follows:

```
<expiration lifespan="1000" maxIdle="1000" />
```

- In Remote Client Server mode, set the ***max-idle*** as follows:

```
<expiration lifespan="1000" max-idle="1000" />
```

### Result

Expiration is now configured for the cache implementation.

[Report a bug](#)

## 4.6. MORTAL AND IMMORTAL DATA

In addition to storing entities, Red Hat JBoss Data Grid allows you to attach mortality information to data. For example, using the standard **put(key, value)** creates an entry that will never expire, called an immortal entry. Alternatively, an entry created using **put(key, value, lifespan, timeunit)** is a mortal entry that has a specified fixed life span, after which it expires.

In addition to the **lifespan** parameter, JBoss Data Grid also provides a **maxIdle** parameter used to determine expiration. The **maxIdle** and **lifespan** parameters can be used in various combinations to set the life span of an entry.

As a default, newly created entries do not have a life span or maximum idle time value set. Without these two values, a data entry will never expire and is therefore known as immortal data.

Set entry mortality (or its expiration values) by setting the life span and maximum idle time values for the entry. After being set, these values must be persisted in the cache stores to ensure that they survive eviction and passivation.

[Report a bug](#)

## 4.7. TROUBLESHOOTING EXPIRATION

If expiration does not appear to be working, it may be due to an entry being marked for expiration but not being removed.

Multiple-cache operations such as **put()** are passed a life span value as a parameter. This value defines the interval after which the entry must expire. In cases where eviction is not configured and the life span interval expires, it can appear as if Red Hat JBoss Data Grid has not removed the entry. For example, when viewing JMX statistics, such as the **number of entries**, you may see an out of date count, or the persistent store associated with JBoss Data Grid may still contain this entry. Behind the scenes, JBoss Data Grid has marked it as an expired entry, but has not removed it. Removal of such entries happens as follows:

- Enabling the eviction feature causes the eviction thread to periodically detect and purge expired entries.

Any attempt to use **get()** or **containsKey()** for the expired entry causes JBoss Data Grid to return a null value. The expired entry is later removed by the eviction thread.

[Report a bug](#)

# PART III. MONITOR YOUR CACHE

## CHAPTER 5. SET UP LOGGING

### 5.1. ABOUT LOGGING

Red Hat JBoss Data Grid provides highly configurable logging facilities for both its own internal use and for use by deployed applications. The logging subsystem is based on JBoss LogManager and it supports several third party application logging frameworks in addition to JBoss Logging.

The logging subsystem is configured using a system of log categories and log handlers. Log categories define what messages to capture, and log handlers define how to deal with those messages (write to disk, send to console, etc).

After a JBoss Data Grid cache is configured with operations such as eviction and expiration, logging tracks relevant activity (including errors or failures).

When set up correctly, logging provides a detailed account at what occurred in the environment and when. Logging also helps track activity that occurred just before a crash or problem in the environment. This information is useful when troubleshooting or when attempting to identify the source of a crash or error.

[Report a bug](#)

### 5.2. SUPPORTED APPLICATION LOGGING FRAMEWORKS

Red Hat JBoss LogManager supports the following logging frameworks:

- JBoss Logging, which is included with Red Hat JBoss Data Grid 6.
- [Apache Commons Logging](#)
- [Simple Logging Facade for Java \(SLF4J\)](#)
- [Apache log4j](#)
- [Java SE Logging \(java.util.logging\)](#)

[Report a bug](#)

#### 5.2.1. About JBoss Logging

JBoss Logging is the application logging framework that is included in JBoss Enterprise Application Platform 6. As a result of this inclusion, Red Hat JBoss Data Grid 6 also uses JBoss Logging.

JBoss Logging provides an easy way to add logging to an application. Add code to the application that uses the framework to send log messages in a defined format. When the application is deployed to an application server, these messages can be captured by the server and displayed and/or written to file according to the server's configuration.

[Report a bug](#)

#### 5.2.2. JBoss Logging Features

JBoss Logging includes the following features:

- Provides an innovative, easy to use *typed* logger.
- Full support for internationalization and localization. Translators work with message bundles in properties files while developers can work with interfaces and annotations.
- Build-time tooling to generate typed loggers for production, and runtime generation of typed loggers for development.

[Report a bug](#)

## 5.3. BOOT LOGGING

The boot log is the record of events that occur while the server is starting up (or booting). Red Hat JBoss Data Grid also includes a server log, which includes log entries generated after the server concludes the boot process.

[Report a bug](#)

### 5.3.1. Configure Boot Logging

Edit the **logging.properties** file to configure the boot log. This file is a standard Java properties file and can be edited in a text editor. Each line in the file has the format of **property=value**.

In Red Hat JBoss Data Grid, the **logging.properties** file is available in the **\$JDG\_HOME/standalone/configuration** folder.

[Report a bug](#)

### 5.3.2. Default Log File Locations

The following table provides a list of log files in Red Hat JBoss Data Grid and their locations:

**Table 5.1. Default Log File Locations**

Log File	Location	Description
<b>boot.log</b>	<b>\$JDG_HOME/standalone/log/</b>	The Server Boot Log. Contains log messages related to the start up of the server.
<b>server.log</b>	<b>\$JDG_HOME/standalone/log/</b>	The Server Log. Contains all log messages once the server has launched.

[Report a bug](#)

## 5.4. LOGGING ATTRIBUTES

### 5.4.1. About Log Levels

Log levels are an ordered set of enumerated values that indicate the nature and severity of a log message. The level of a given log message is specified by the developer using the appropriate methods of their chosen logging framework to send the message.

Red Hat JBoss Data Grid supports all the log levels used by the supported application logging frameworks. The six most commonly used log levels are (ordered by lowest to highest severity):

1. **TRACE**
2. **DEBUG**
3. **INFO**
4. **WARN**
5. **ERROR**
6. **FATAL**

Log levels are used by log categories and handlers to limit the messages they are responsible for. Each log level has an assigned numeric value which indicates its order relative to other log levels. Log categories and handlers are assigned a log level and they only process log messages of that numeric value or higher. For example a log handler with the level of **WARN** will only record messages of the levels **WARN**, **ERROR** and **FATAL**.

[Report a bug](#)

## 5.4.2. Supported Log Levels

The following table lists log levels that are supported in Red Hat JBoss Data Grid. Each entry includes the log level, its value and description. The log level values indicate each log level's relative value to other log levels. Additionally, log levels in different frameworks may be named differently, but have a log value consistent to the provided list.

**Table 5.2. Supported Log Levels**

Log Level	Value	Description
FINEST	300	-
FINER	400	-
TRACE	400	Used for messages that provide detailed information about the running state of an application. <b>TRACE</b> level log messages are captured when the server runs with the <b>TRACE</b> level enabled.

Log Level	Value	Description
DEBUG	500	Used for messages that indicate the progress of individual requests or activities of an application. <b>DEBUG</b> level log messages are captured when the server runs with the <b>DEBUG</b> level enabled.
FINE	500	-
CONFIG	700	-
INFO	800	Used for messages that indicate the overall progress of the application. Used for application start up, shut down and other major lifecycle events.
WARN	900	Used to indicate a situation that is not in error but is not considered ideal. Indicates circumstances that can lead to errors in the future.
WARNING	900	-
ERROR	1000	Used to indicate an error that has occurred that could prevent the current activity or request from completing but will not prevent the application from running.
SEVERE	1000	-
FATAL	1100	Used to indicate events that could cause critical service failure and application shutdown and possibly cause JBoss Data Grid to shut down.

[Report a bug](#)

### 5.4.3. About Log Categories

Log categories define a set of log messages to capture and one or more log handlers which will process the messages.

The log messages to capture are defined by their Java package of origin and log level. Messages from classes in that package and of that log level or higher (with greater or equal numeric value) are captured by the log category and sent to the specified log handlers. As an example, the **WARNING** log level results in log values of **900**, **1000** and **1100** are captured.

Log categories can optionally use the log handlers of the root logger instead of their own handlers.

[Report a bug](#)

#### 5.4.4. About the Root Logger

The root logger captures all log messages sent to the server (of a specified level) that are not captured by a log category. These messages are then sent to one or more log handlers.

By default the root logger is configured to use a console and a periodic log handler. The periodic log handler is configured to write to the file `server.log`. This file is sometimes referred to as the server log.

[Report a bug](#)

#### 5.4.5. About Log Handlers

Log handlers define how captured log messages are recorded by Red Hat JBoss Data Grid. The six types of log handlers configurable in JBoss Data Grid are:

- **Console**
- **File**
- **Periodic**
- **Size**
- **Async**
- **Custom**

Log handlers direct specified log objects to a variety of outputs (including the console or specified log files). Some log handlers used in JBoss Data Grid are wrapper log handlers, used to direct other log handlers' behavior.

Log handlers are used to direct log outputs to specific files for easier sorting or to write logs for specific intervals of time. They are primarily useful to specify the kind of logs required and where they are stored or displayed or the logging behavior in JBoss Data Grid.

[Report a bug](#)

#### 5.4.6. Log Handler Types

The following table lists the different types of log handlers available in Red Hat JBoss Data Grid:

**Table 5.3. Log Handler Types**

Log Handler Type	Description	Use Case
------------------	-------------	----------



Log Handler Type	Description	Use Case
Console	Console log handlers write log messages to either the host operating system's standard out ( <b>stdout</b> ) or standard error ( <b>stderr</b> ) stream. These messages are displayed when JBoss Data Grid is run from a command line prompt.	The Console log handler is preferred when JBoss Data Grid is administered using the command line. In such a case, the messages from a Console log handler are not saved unless the operating system is configured to capture the standard out or standard error stream.
File	File log handlers are the simplest log handlers. Their primary use is to write log messages to a specified file.	File log handlers are most useful if the requirement is to store all log entries according to the time in one place.
Periodic	Periodic file handlers write log messages to a named file until a specified period of time has elapsed. Once the time period has elapsed, the specified time stamp is appended to the file name. The handler then continues to write into the newly created log file with the original name.	The Periodic file handler can be used to accumulate log messages on a weekly, daily, hourly or other basis depending on the requirements of the environment.
Size	Size log handlers write log messages to a named file until the file reaches a specified size. When the file reaches a specified size, it is renamed with a numeric prefix and the handler continues to write into a newly created log file with the original name. Each size log handler must specify the maximum number of files to be kept in this fashion.	The Size handler is best suited to an environment where the log file size must be consistent.
Async	Async log handlers are wrapper log handlers that provide asynchronous behavior for one or more other log handlers. These are useful for log handlers that have high latency or other performance problems such as writing a log file to a network file system.	The Async log handlers are best suited to an environment where high latency is a problem or when writing to a network file system.

Log Handler Type	Description	Use Case
Custom	Custom log handlers enable you to configure new types of log handlers that have been implemented. A custom handler must be implemented as a Java class that extends <b>java.util.logging.Handler</b> and be contained in a module.	Custom log handlers create customized log handler types and are recommended for advanced users.

[Report a bug](#)

### 5.4.7. Selecting Log Handlers

The following are the most common uses for each of the log handler types available for Red Hat JBoss Data Grid:

- The **Console** log handler is preferred when JBoss Data Grid is administered using the command line. In such a case, errors and log messages appear on the console window and are not saved unless separately configured to do so.
- The **File** log handler is used to direct log entries into a specified file. This simplicity is useful if the requirement is to store all log entries according to the time in one place.
- The **Periodic** log handler is similar to the **File** handler but creates files according to the specified period. As an example, this handler can be used to accumulate log messages on a weekly, daily, hourly or other basis depending on the requirements of the environment.
- The **Size** log handler also writes log messages to a specified file, but only while the log file size is within a specified limit. Once the file size reaches the specified limit, log files are written to a new log file. This handler is best suited to an environment where the log file size must be consistent.
- The **Async** log handler is a wrapper that forces other log handlers to operate asynchronously. This is best suited to an environment where high latency is a problem or when writing to a network file system.
- The **Custom** log handler creates new, customized types of log handlers. This is an advanced log handler.

[Report a bug](#)

### 5.4.8. About Log Formatters

A log formatter is the configuration property of a log handler. The log formatter defines the appearance of log messages that originate from the relevant log handler. The log formatter is a string that uses the same syntax as the **java.util.Formatter** class.

See <http://docs.oracle.com/javase/6/docs/api/java/util/Formatter.html> for more information.

[Report a bug](#)

## 5.5. LOGGING SAMPLE CONFIGURATIONS

### 5.5.1. Sample XML Configuration for the Root Logger

The following procedure demonstrates a sample configuration for the root logger.

#### Procedure 5.1. Configure the Root Logger

##### 1. Set the *level* Property

The *level* property sets the maximum level of log message that the root logger records.

```
<subsystem xmlns="urn:jboss:domain:logging:1.4">
  <root-logger>
    <level name="INFO"/>
```

##### 2. List *handlers*

*handlers* is a list of log handlers that are used by the root logger.

```
<subsystem xmlns="urn:jboss:domain:logging:1.4">
  <root-logger>
    <level name="INFO"/>
    <handlers>
      <handler name="CONSOLE"/>
      <handler name="FILE"/>
    </handlers>
  </root-logger>
</subsystem>
```

[Report a bug](#)

### 5.5.2. Sample XML Configuration for a Log Category

The following procedure demonstrates a sample configuration for a log category.

#### Procedure 5.2. Configure a Log Category

```
<subsystem xmlns="urn:jboss:domain:logging:1.4">
  <logger category="com.company.accounts.rec" use-parent-handlers="true">
    <level name="WARN"/>
    <handlers>
      <handler name="accounts-rec"/>
    </handlers>
  </logger>
</subsystem>
```

1. Use the *category* property to specify the log category from which log messages will be captured.

The *use-parent-handlers* is set to **"true"** by default. When set to **"true"**, this category will use the log handlers of the root logger in addition to any other assigned handlers.

2. Use the *level* property to set the maximum level of log message that the log category records.

3. The **handlers** element contains a list of log handlers.

[Report a bug](#)

### 5.5.3. Sample XML Configuration for a Console Log Handler

The following procedure demonstrates a sample configuration for a console log handler.

#### Procedure 5.3. Configure the Console Log Handler

```
<subsystem xmlns="urn:jboss:domain:logging:1.4">
  <console-handler name="CONSOLE" autoflush="true">
    <level name="INFO"/>
    <encoding value="UTF-8"/>
    <target value="System.out"/>
    <filter-spec value="not(match(&quot;JBAS.*&quot;))"/>
    <formatter>
      <pattern-formatter pattern="%K{level}%d{HH:mm:ss,SSS} %-5p [%c]
(%t) %s%E%n"/>
    </formatter>
  </console-handler>
</subsystem>
```

1. **Add the Log Handler Identifier Information**

The **name** property sets the unique identifier for this log handler.

When **autoflush** is set to **"true"** the log messages will be sent to the handler's target immediately upon request.

2. **Set the *level* Property**

The **level** property sets the maximum level of log messages recorded.

3. **Set the *encoding* Output**

Use **encoding** to set the character encoding scheme to be used for the output.

4. **Define the *target* Value**

The **target** property defines the system output stream where the output of the log handler goes. This can be **System.err** for the system error stream, or **System.out** for the standard out stream.

5. **Define the *filter-spec* Property**

The **filter-spec** property is an expression value that defines a filter. The example provided defines a filter that does not match a pattern: **not(match("JBAS.\*"))**.

6. **Specify the *formatter***

Use **formatter** to list the log formatter used by the log handler.

[Report a bug](#)

### 5.5.4. Sample XML Configuration for a File Log Handler

The following procedure demonstrates a sample configuration for a file log handler.

## Procedure 5.4. Configure the File Log Handler

```
<file-handler name="accounts-rec-trail" autoflush="true">
  <level name="INFO"/>
  <encoding value="UTF-8"/>
  <file relative-to="jboss.server.log.dir" path="accounts-rec-
trail.log"/>
  <formatter>
    <pattern-formatter pattern="%d{HH:mm:ss,SSS} %-5p [%c] (%t)
%S%E%n"/>
  </formatter>
  <append value="true"/>
</file-handler>
```

### 1. Add the File Log Handler Identifier Information

The *name* property sets the unique identifier for this log handler.

When *autoflush* is set to **"true"** the log messages will be sent to the handler's target immediately upon request.

### 2. Set the *level* Property

The *level* property sets the maximum level of log message that the root logger records.

### 3. Set the *encoding* Output

Use *encoding* to set the character encoding scheme to be used for the output.

### 4. Set the *file* Object

The *file* object represents the file where the output of this log handler is written to. It has two configuration properties: *relative-to* and *path*.

The *relative-to* property is the directory where the log file is written to. JBoss Enterprise Application Platform 6 file path variables can be specified here. The `jboss.server.log.dir` variable points to the `log/` directory of the server.

The *path* property is the name of the file where the log messages will be written. It is a relative path name that is appended to the value of the *relative-to* property to determine the complete path.

### 5. Specify the *formatter*

Use *formatter* to list the log formatter used by the log handler.

### 6. Set the *append* Property

When the *append* property is set to **"true"**, all messages written by this handler will be appended to an existing file. If set to **"false"** a new file will be created each time the application server launches. Changes to *append* require a server reboot to take effect.

[Report a bug](#)

## 5.5.5. Sample XML Configuration for a Periodic Log Handler

The following procedure demonstrates a sample configuration for a periodic log handler.

### Procedure 5.5. Configure the Periodic Log Handler

```

<periodic-rotating-file-handler name="FILE" autoflush="true">
  <level name="INFO"/>
  <encoding value="UTF-8"/>
  <formatter>
    <pattern-formatter pattern="%d{HH:mm:ss,SSS} %-5p [%c] (%t)
%s%E%n"/>
  </formatter>
  <file relative-to="jboss.server.log.dir" path="server.log"/>
  <suffix value=".yyyy-MM-dd"/>
  <append value="true"/>
</periodic-rotating-file-handler>

```

### 1. Add the Periodic Log Handler Identifier Information

The **name** property sets the unique identifier for this log handler.

When **autoflush** is set to **"true"** the log messages will be sent to the handler's target immediately upon request.

### 2. Set the *level* Property

The **level** property sets the maximum level of log message that the root logger records.

### 3. Set the *encoding* Output

Use **encoding** to set the character encoding scheme to be used for the output.

### 4. Specify the *formatter*

Use **formatter** to list the log formatter used by the log handler.

### 5. Set the *file* Object

The **file** object represents the file where the output of this log handler is written to. It has two configuration properties: **relative-to** and **path**.

The **relative-to** property is the directory where the log file is written to. JBoss Enterprise Application Platform 6 file path variables can be specified here. The **jboss.server.log.dir** variable points to the **log/** directory of the server.

The **path** property is the name of the file where the log messages will be written. It is a relative path name that is appended to the value of the **relative-to** property to determine the complete path.

### 6. Set the *suffix* Value

The **suffix** is appended to the filename of the rotated logs and is used to determine the frequency of rotation. The format of the **suffix** is a dot (.) followed by a date string, which is parsable by the **java.text.SimpleDateFormat** class. The log is rotated on the basis of the smallest time unit defined by the **suffix**. For example, **yyyy-MM-dd** will result in daily log rotation. See <http://docs.oracle.com/javase/6/docs/api/index.html?java/text/SimpleDateFormat.html>

### 7. Set the *append* Property

When the **append** property is set to **"true"**, all messages written by this handler will be appended to an existing file. If set to **"false"** a new file will be created each time the application server launches. Changes to **append** require a server reboot to take effect.

[Report a bug](#)

## 5.5.6. Sample XML Configuration for a Size Log Handler

The following procedure demonstrates a sample configuration for a size log handler.

### Procedure 5.6. Configure the Size Log Handler

```
<size-rotating-file-handler name="accounts_debug" autoflush="false">
  <level name="DEBUG"/>
  <encoding value="UTF-8"/>
  <file relative-to="jboss.server.log.dir" path="accounts-debug.log"/>
  <rotate-size value="500k"/>
  <max-backup-index value="5"/>
  <formatter>
    <pattern-formatter pattern="%d{HH:mm:ss,SSS} %-5p [%c] (%t)
%S%E%n"/>
  </formatter>
  <append value="true"/>
</size-rotating-file-handler>
```

#### 1. Add the Size Log Handler Identifier Information

The *name* property sets the unique identifier for this log handler.

When *autoflush* is set to "true" the log messages will be sent to the handler's target immediately upon request.

#### 2. Set the *level* Property

The *level* property sets the maximum level of log message that the root logger records.

#### 3. Set the *encoding* Output

Use *encoding* to set the character encoding scheme to be used for the output.

#### 4. Set the *file* Object

The *file* object represents the file where the output of this log handler is written to. It has two configuration properties: *relative-to* and *path*.

The *relative-to* property is the directory where the log file is written to. JBoss Enterprise Application Platform 6 file path variables can be specified here. The `jboss.server.log.dir` variable points to the `log/` directory of the server.

The *path* property is the name of the file where the log messages will be written. It is a relative path name that is appended to the value of the *relative-to* property to determine the complete path.

#### 5. Specify the *rotate-size* Value

The maximum size that the log file can reach before it is rotated. A single character appended to the number indicates the size units: **b** for bytes, **k** for kilobytes, **m** for megabytes, **g** for gigabytes. For example: **50m** for 50 megabytes.

#### 6. Set the *max-backup-index* Number

The maximum number of rotated logs that are kept. When this number is reached, the oldest log is reused.

#### 7. Specify the *formatter*

Use *formatter* to list the log formatter used by the log handler.

### 8. Set the *append* Property

When the *append* property is set to "**true**", all messages written by this handler will be appended to an existing file. If set to "**false**" a new file will be created each time the application server launches. Changes to *append* require a server reboot to take effect.

[Report a bug](#)

## 5.5.7. Sample XML Configuration for a Async Log Handler

The following procedure demonstrates a sample configuration for an async log handler

### Procedure 5.7. Configure the Async Log Handler

```
<async-handler name="Async_NFS_handlers">
  <level name="INFO"/>
  <queue-length value="512"/>
  <overflow-action value="block"/>
  <subhandlers>
    <handler name="FILE"/>
    <handler name="accounts-record"/>
  </subhandlers>
</async-handler>
```

1. The *name* property sets the unique identifier for this log handler.
2. The *level* property sets the maximum level of log message that the root logger records.
3. The *queue-length* defines the maximum number of log messages that will be held by this handler while waiting for sub-handlers to respond.
4. The *overflow-action* defines how this handler responds when its queue length is exceeded. This can be set to **BLOCK** or **DISCARD**. **BLOCK** makes the logging application wait until there is available space in the queue. This is the same behavior as an non-async log handler. **DISCARD** allows the logging application to continue but the log message is deleted.
5. The *subhandlers* list is the list of log handlers to which this async handler passes its log messages.

[Report a bug](#)



## PART IV. SET UP CACHE MODES

## CHAPTER 6. CACHE MODES

Red Hat JBoss Data Grid provides two modes:

- Local mode is the only non-clustered cache mode offered in JBoss Data Grid. In local mode, JBoss Data Grid operates as a simple single-node in-memory data cache. Local mode is most effective when scalability and failover are not required and provides high performance in comparison with clustered modes.
- Clustered mode replicates state changes to a small subset of nodes. The subset size is sufficient for fault tolerance purposes but not large enough to hinder scalability. Before attempting to use clustered mode, it is important to first configure JGroups for a clustered configuration. For details about configuring JGroups, see [Section 26.2, “Configure JGroups \(Library Mode\)”](#)

[Report a bug](#)

### 6.1. ABOUT CACHE CONTAINERS

Cache containers are used in Red Hat JBoss Data Grid's Remote Client-Server mode as a starting point for a cache. The **cache-container** element acts as a parent of one or more (local or clustered) caches. To add clustered caches to the container, transport must be defined.

The following procedure demonstrates a sample cache container configuration:

#### Procedure 6.1. How to Configure the Cache Container

```
<subsystem xmlns="urn:infinispan:server:core:6.2"
  default-cache-container="local">
  <cache-container name="local"
    default-cache="default"
    statistics="true"
    listener-executor="infinispan-listener"
    start="EAGER">
    <local-cache name="default"
      start="EAGER"
      statistics="false">
      <!-- Additional configuration information here -->
    </local-cache>
  </cache-container>
</subsystem>
```

#### 1. Configure the Cache Container

The **cache-container** element specifies information about the cache container using the following parameters:

- a. The **name** parameter defines the name of the cache container.
- b. The **default-cache** parameter defines the name of the default cache used with the cache container.
- c. The **statistics** attribute is optional and is **true** by default. Statistics are useful in monitoring JBoss Data Grid via JMX or JBoss Operations Network, however they adversely affect performance. Disable this attribute by setting it to **false** if it is not required.

- d. The *listener-executor* defines the executor used for asynchronous cache listener notifications.
- e. The *start* parameter indicates when the cache container starts, i.e. whether it will start lazily when requested or "eagerly" when the server starts up. Valid values for this parameter are **EAGER** and **LAZY**.

## 2. Configure Per-cache Statistics

If *statistics* are enabled at the container level, per-cache statistics can be selectively disabled for caches that do not require monitoring by setting the *statistics* attribute to **false**.

[Report a bug](#)

## 6.2. LOCAL MODE

Using Red Hat JBoss Data Grid's local mode instead of a map provides a number of benefits.

Caches offer features that are unmatched by simple maps, such as:

- Write-through and write-behind caching to persist data.
- Entry eviction to prevent the Java Virtual Machine (JVM) running out of memory.
- Support for entries that expire after a defined period.

JBoss Data Grid is built around a high performance, read-based data container that uses techniques such as optimistic and pessimistic locking to manage lock acquisitions.

JBoss Data Grid also uses compare-and-swap and other lock-free algorithms, resulting in high throughput multi-CPU or multi-core environments. Additionally, JBoss Data Grid's Cache API extends the JDK's **ConcurrentMap**, resulting in a simple migration process from a map to JBoss Data Grid.

[Report a bug](#)

### 6.2.1. Configure Local Mode (Remote Client-Server Mode)

A local cache can be added to any cache container. The following example demonstrates how to add the **local-cache** element.

#### Procedure 6.2. The local-cache Element

```
<cache-container name="local"
    default-cache="default"
    statistics="true">
  <local-cache name="default"
    start="EAGER"
    batching="false"
    statistics="true">
    <indexing index="NONE">
    <property
name="default.directory_provider">ram</property>
    </indexing>
  </local-cache>
```

The **local-cache** element specifies information about the local cache used with the cache container using the following parameters:

1. The **name** parameter specifies the name of the local cache to use.
2. The **start** parameter indicates when the cache container starts, i.e. whether it will start lazily when requested or "eagerly" when the server starts up. Valid values for this parameter are **EAGER** and **LAZY**.
3. The **batching** parameter specifies whether batching is enabled for the local cache.
4. If **statistics** are enabled at the container level, per-cache statistics can be selectively disabled for caches that do not require monitoring by setting the **statistics** attribute to **false**.
5. The **indexing** parameter specifies the type of indexing used for the local cache. Valid values for this parameter are **NONE**, **LOCAL** and **ALL**.

Alternatively, create a **DefaultCacheManager** with the "no-argument" constructor. Both of these methods create a local default cache.

Local and clustered caches are able to coexist in the same cache container, however where the container is without a `<transport/>` it can only contain local caches. The container used in the example can only contain local caches as it does not have a `<transport/>`.

The cache interface extends the **ConcurrentMap** and is compatible with multiple cache systems.

[Report a bug](#)

### 6.2.2. Configure Local Mode (Library Mode)

In Red Hat JBoss Data Grid's Library mode, setting a cache's **mode** parameter to **local** equals not specifying a clustering mode at all. In the case of the latter, the cache defaults to local mode, even if its cache manager defines a transport.

Set the cluster mode to local as follows:

```
<clustering mode="local" />
```

[Report a bug](#)

## 6.3. CLUSTERED MODES

Red Hat JBoss Data Grid offers the following clustered modes:

- Replication Mode replicates any entry that is added across all cache instances in the cluster.
- Invalidation Mode does not share any data, but signals remote caches to initiate the removal of invalid entries.
- Distribution Mode stores each entry on a subset of nodes instead of on all nodes in the cluster.

The clustered modes can be further configured to use synchronous or asynchronous transport for network communications.

[Report a bug](#)

### 6.3.1. Asynchronous and Synchronous Operations

When a clustered mode (such as invalidation, replication or distribution) is used, data is propagated to other nodes in either a synchronous or asynchronous manner.

If synchronous mode is used, the sender waits for responses from receivers before allowing the thread to continue, whereas asynchronous mode transmits data but does not wait for responses from other nodes in the cluster to continue operations.

Asynchronous mode prioritizes speed over consistency, which is ideal for use cases such as HTTP session replications with sticky sessions enabled. Such a session (or data for other use cases) is always accessed on the same cluster node, unless this node fails.

[Report a bug](#)

### 6.3.2. Cache Mode Troubleshooting

#### 6.3.2.1. Invalid Data in ReadExternal

If invalid data is passed to `readExternal`, it can be because when using `Cache.putAsync()`, starting serialization can cause your object to be modified, causing the datastream passed to `readExternal` to be corrupted. This can be resolved if access to the object is synchronized.

[Report a bug](#)

#### 6.3.2.2. About Asynchronous Communications

In Red Hat JBoss Data Grid, the local, distributed and replicated modes are represented by the `local-cache`, `distributed-cache` and `replicated-cache` elements respectively. Each of these elements contains a `mode` property, the value of which can be set to `SYNC` for synchronous or `ASYNC` for asynchronous communications.

#### Example 6.1. Asynchronous Communications Example Configuration

```
<replicated-cache name="default"
    start="EAGER"
    mode="SYNC"
    batching="false"
    statistics="true">
    <!-- Additional configuration information here -->
</replicated-cache>
```



#### NOTE

This configuration is valid for both JBoss Data Grid's usage modes (Library mode and Remote Client-Server mode).

[Report a bug](#)

### 6.3.2.3. Cluster Physical Address Retrieval

#### How can the physical addresses of the cluster be retrieved?

The physical address can be retrieved using an instance method call. For example:

```
AdvancedCache.getRpcManager().getTransport().getPhysicalAddresses().
```

[Report a bug](#)

## 6.4. STATE TRANSFER

State transfer is a basic data grid or clustered cache functionality. Without state transfer, data would be lost as nodes are added to or removed from the cluster.

State transfer adjusts the cache's internal state in response to a change in a cache membership. The change can be when a node joins or leaves, when two or more cluster partitions merge, or a combination of joins, leaves, and merges. State transfer occurs automatically in Red Hat JBoss Data Grid whenever a node joins or leaves the cluster.

In Red Hat JBoss Data Grid's replication mode, a new node joining the cache receives the entire cache state from the existing nodes. In distribution mode, the new node receives only a part of the state from the existing nodes, and the existing nodes remove some of their state in order to keep **numOwners** copies of each key in the cache (as determined through consistent hashing). In invalidation mode the initial state transfer is similar to replication mode, the only difference being that the nodes are not guaranteed to have the same state. When a node leaves, a replicated mode or invalidation mode cache does not perform any state transfer. A distributed cache needs to make additional copies of the keys that were stored on the leaving nodes, again to keep **numOwners** copies of each key.

A State Transfer transfers both in-memory and persistent state by default, but both can be disabled in the configuration. When State Transfer is disabled a **ClusterLoader** must be configured, otherwise a node will become the owner or backup owner of a key without the data being loaded into its cache. In addition, if State Transfer is disabled in distributed mode then a key will occasionally have less than **numOwners** owners.

[Report a bug](#)

### 6.4.1. Non-Blocking State Transfer

Non-Blocking State Transfer in Red Hat JBoss Data Grid minimizes the time in which a cluster or node is unable to respond due to a state transfer in progress. Non-blocking state transfer is a core architectural improvement with the following goals:

- Minimize the interval(s) where the entire cluster cannot respond to requests because of a state transfer in progress.
- Minimize the interval(s) where an existing member stops responding to requests because of a state transfer in progress.
- Allow state transfer to occur with a drop in the performance of the cluster. However, the drop in the performance during the state transfer does not throw any exception, and allows processes to continue.
- Allows a **GET** operation to successfully retrieve a key from another node without returning a null value during a progressive state transfer.

For simplicity, the total order-based commit protocol uses a blocking version of the currently implemented state transfer mechanism. The main differences between the regular state transfer and the total order state transfer are:

- The blocking protocol queues the transaction delivery during the state transfer.
- State transfer control messages (such as `CacheTopologyControlCommand`) are sent according to the total order information.

The total order-based commit protocol works with the assumption that all the transactions are delivered in the same order and they see the same data set. So, no transactions are validated during the state transfer because all the nodes must have the most recent key or values in memory.

Using the state transfer and blocking protocol in this manner allows the state transfer and transaction delivery on all on the nodes to be synchronized. However, transactions that are already involved in a state transfer (sent before the state transfer began and delivered after it concludes) must be resent. When resent, these transactions are treated as new joiners and assigned a new total order value.

[Report a bug](#)

### 6.4.2. Suppress State Transfer via JMX

State transfer can be suppressed using JMX in order to bring down and relaunch a cluster for maintenance. This operation permits a more efficient cluster shutdown and startup, and removes the risk of Out Of Memory errors when bringing down a grid.

When a new node joins the cluster and rebalancing is suspended, the `getCache()` call will timeout after `stateTransfer.timeout` expires unless rebalancing is re-enabled or `stateTransfer.awaitInitialTransfer` is set to `false`.

Disabling state transfer and rebalancing can be used for partial cluster shutdown or restart, however there is the possibility that data may be lost in a partial cluster shutdown due to state transfer being disabled.

[Report a bug](#)

### 6.4.3. The `rebalancingEnabled` Attribute

Suppressing rebalancing can only be triggered via the `rebalancingEnabled` JMX attribute, and requires no specific configuration.

The `rebalancingEnabled` attribute can be modified for the entire cluster from the `LocalTopologyManager` JMX Mbean on any node. This attribute is `true` by default, and is configurable programmatically.

Servers such as Hot Rod attempt to start all caches declared in the configuration during startup. If rebalancing is disabled, the cache will fail to start. Therefore, it is mandatory to use the following setting in a server environment:

```
<await-initial-transfer="false"/>
```

[Report a bug](#)

## CHAPTER 7. SET UP DISTRIBUTION MODE

### 7.1. ABOUT DISTRIBUTION MODE

When enabled, Red Hat JBoss Data Grid's distribution mode stores each entry on a subset of the nodes in the grid instead of replicating each entry on every node. Typically, each entry is stored on more than one node for redundancy and fault tolerance.

As a result of storing entries on selected nodes across the cluster, distribution mode provides improved scalability compared to other clustered modes.

A cache using distribution mode can transparently locate keys across a cluster using the consistent hash algorithm.

[Report a bug](#)

### 7.2. DISTRIBUTION MODE'S CONSISTENT HASH ALGORITHM

The hashing algorithm in Red Hat JBoss Data Grid is based on consistent hashing. The term consistent hashing is still used for this implementation, despite some divergence from a traditional consistent hash.

Distribution mode uses a consistent hash algorithm to select a node from the cluster to store entries upon. The consistent hash algorithm is configured with the number of copies of each cache entry to be maintained within the cluster. Unlike generic consistent hashing, the implementation used in JBoss Data Grid splits the key space into fixed segments. The number of segments is configurable using *numSegments* and cannot be changed without restarting the cluster. The mapping of keys to segments is also fixed — a key maps to the same segment, regardless of how the topology of the cluster changes.

The number of copies set for each data item requires balancing performance and fault tolerance. Creating too many copies of the entry can impair performance and too few copies can result in data loss in case of node failure.

Each hash segment is mapped to a list of nodes called owners. The order is important because the first owner (also known as the primary owner) has a special role in many cache operations (for example, locking). The other owners are called backup owners. There is no rule about mapping segments to owners, although the hashing algorithms simultaneously balance the number of segments allocated to each node and minimize the number of segments that have to move after a node joins or leaves the cluster.

[Report a bug](#)

### 7.3. LOCATING ENTRIES IN DISTRIBUTION MODE

The consistent hash algorithm used in Red Hat JBoss Data Grid's distribution mode can locate entries deterministically, without multicasting a request or maintaining expensive metadata.

A **PUT** operation can result in as many remote calls as specified by the *num\_copies* parameter, while a **GET** operation executed on any node in the cluster results in a single remote call. In the background, the **GET** operation results in the same number of remote calls as a **PUT** operation (specifically the value of the *num\_copies* parameter), but these occur in parallel and the returned entry is passed to the caller as soon as one returns.

[Report a bug](#)



## 7.4. RETURN VALUES IN DISTRIBUTION MODE

In Red Hat JBoss Data Grid's distribution mode, a synchronous request is used to retrieve the previous return value if it cannot be found locally. A synchronous request is used for this task irrespective of whether distribution mode is using asynchronous or synchronous processes.

[Report a bug](#)

## 7.5. CONFIGURE DISTRIBUTION MODE (REMOTE CLIENT-SERVER MODE)

Distribution mode is a clustered mode in Red Hat JBoss Data Grid. Distribution mode can be added to any cache container using the following procedure:

### Procedure 7.1. The `distributed-cache` Element

```
<cache-container name="clustered"
  default-cache="default"
  statistics="true">
  <transport executor="infinispan-transport" lock-timeout="60000"/>
  <distributed-cache name="default"
    mode="SYNC"
    segments="20"
    start="EAGER"
    statistics="true">
    <!-- Additional configuration information here -->
  </distributed-cache>
</cache-container>
```

The `distributed-cache` element configures settings for the distributed cache using the following parameters:

1. The **name** parameter provides a unique identifier for the cache.
2. The **mode** parameter sets the clustered cache mode. Valid values are **SYNC** (synchronous) and **ASYNC** (asynchronous).
3. The (optional) **segments** parameter specifies the number of hash space segments per cluster. The recommended value for this parameter is ten multiplied by the cluster size and the default value is **20**.
4. The **start** parameter specifies whether the cache starts when the server starts up or when it is requested or deployed.
5. If **statistics** are enabled at the container level, per-cache statistics can be selectively disabled for caches that do not require monitoring by setting the **statistics** attribute to **false**.



### IMPORTANT

JGroups must be appropriately configured for clustered mode before attempting to load this configuration.

For details about the **cache-container**, **locking**, and **transaction** elements, see the appropriate chapter.

[Report a bug](#)

## 7.6. CONFIGURE DISTRIBUTION MODE (LIBRARY MODE)

The following procedure shows a distributed cache configuration in Red Hat JBoss Data Grid's Library mode.

### Procedure 7.2. Distributed Cache Configuration

#### 1. Configure the Clustering Element

```
<clustering mode="distribution">
  <sync replTimeout="${TIME}" />
  <stateTransfer chunkSize="${SIZE}"
    fetchInMemoryState="{true/false}"
    awaitInitialTransfer="{true/false}"
    timeout="${TIME}" />
```

- a. The **clustering** element's **mode** parameter's value determines the clustering mode selected for the cache.
- b. The **sync** element's **replTimeout** parameter specifies the maximum time period in milliseconds for an acknowledgment after a remote call. If the time period ends without any acknowledgment, an exception is thrown.
- c. The **stateTransfer** element specifies how state is transferred when a node leaves or joins the cluster. It uses the following parameters:
  - i. The **chunkSize** parameter is the number of cache entries to be transferred in one chunk. The default **chunkSize** value is 512. The **chunkSize** depends on many parameters (for example, size of entries) and the best setting needs to be measured. To change the value for larger cache entries, smaller chunks are advisable and for smaller cache entries, increasing the **chunkSize** offers better performance.
  - ii. The **fetchInMemoryState** parameter when set to **true**, requests state information from neighboring caches on start up. This impacts the start up time for the cache. Set the default to **true** if there is no shared cache store behind to avoid cache inconsistency.
  - iii. The **awaitInitialTransfer** parameter causes the first call to method **CacheManager.getCache()** on the joiner node to block and wait until the joining is complete and the cache has finished receiving state from neighboring caches (if **fetchInMemoryState** is enabled). This option applies to distributed and replicated caches only and is enabled by default.
  - iv. The **timeout** parameter specifies the maximum time (in milliseconds) the cache waits for responses from neighboring caches with the requested states. If no response is received within the **timeout** period, the start up process aborts and an exception is thrown. As a result of the failed state transfer, the cache is not available on that instance.

#### 2. Specify Transport Configuration

■

```

<global>
  <transport clusterName="${NAME}"
    distributedSyncTimeout="${TIME}"
    transportClass="${CLASS}" />
</global>

```

The **transport** element defines the transport configuration for the cache container as follows:

- The **clusterName** parameter specifies the name of the cluster. Nodes can only connect to clusters that share the same name.
- The **distributedSyncTimeout** parameter specifies the time to wait to acquire a lock on the distributed lock. This distributed lock ensures that a single cache can transfer state or rehash state at a time.
- The **transportClass** parameter specifies a class that represents a network transport for the cache container.

[Report a bug](#)

## 7.7. SYNCHRONOUS AND ASYNCHRONOUS DISTRIBUTION

Distribution mode only supports synchronous communication. To elicit meaningful return values from certain public API methods, it is essential to use synchronized communication when using distribution mode.

### Example 7.1. Communication Mode example

For example, with three caches in a cluster, cache **A**, **B** and **C**, and a key **K** that maps cache **A** to **B**. Perform an operation on cluster **C** that requires a return value, for example **Cache.remove(K)**. To execute successfully, the operation must first synchronously forward the call to both cache **A** and **B**, and then wait for a result returned from either cache **A** or **B**. If asynchronous communication was used, the usefulness of the returned values cannot be guaranteed, despite the operation behaving as expected.

[Report a bug](#)

## 7.8. GET AND PUT USAGE IN DISTRIBUTION MODE

In distribution mode, the cache performs a remote **GET** command before a write command. This occurs because certain methods (for example, **Cache.put()**) return the previous value associated with the specified key according to the **java.util.Map** contract. When this is performed on an instance that does not own the key and the entry is not found in the L1 cache, the only reliable way to elicit this return value is to perform a remote **GET** before the **PUT**.

The **GET** operation that occurs before the **PUT** operation is always synchronous, whether the cache is synchronous or asynchronous, because Red Hat JBoss Data Grid must wait for the return value.

[Report a bug](#)

### 7.8.1. Distributed GET and PUT Operation Resource Usage

In distribution mode, the cache may execute a **GET** operation before executing the desired **PUT** operation.

This operation is very expensive in terms of resources. Despite operating in an synchronous manner, a remote **GET** operation does not wait for all responses, which would result in wasted resources. The **GET** process accepts the first valid response received, which allows its performance to be unrelated to cluster size.

Use the ***Flag.SKIP\_REMOTE\_LOOKUP*** flag for a per-invocation setting if return values are not required for your implementation.

Such actions do not impair cache operations and the accurate functioning of all public methods, but do break the **java.util.Map** interface contract. The contract breaks because unreliable and inaccurate return values are provided to certain methods. As a result, ensure that these return values are not used for any important purpose on your configuration.

[Report a bug](#)

## CHAPTER 8. SET UP REPLICATION MODE

### 8.1. ABOUT REPLICATION MODE

Red Hat JBoss Data Grid's replication mode is a simple clustered mode. Cache instances automatically discover neighboring instances on other Java Virtual Machines (JVM) on the same network and subsequently form a cluster with the discovered instances. Any entry added to a cache instance is replicated across all cache instances in the cluster and can be retrieved locally from any cluster cache instance.

In JBoss Data Grid's replication mode, return values are locally available before the replication occurs.

[Report a bug](#)

### 8.2. OPTIMIZED REPLICATION MODE USAGE

Replication mode is used for state sharing across a cluster; however, if you have a replicated cache and a large number of nodes are in use then there will be many writes to the replicated cache to keep all of the nodes synchronized. The amount of work performed will depend on many factors and on the specific use case, and for this reason it is recommended to ensure that each workload is tested thoroughly to determine if replication mode will be beneficial with the number of planned nodes. For many situations replication mode is not recommended once there are ten servers; however, in some workloads, such as if load read is important, this mode may be beneficial.

Red Hat JBoss Data Grid can be configured to use UDP multicast, which improves performance to a limited degree for larger clusters.


[Report a bug](#)

### 8.3. CONFIGURE REPLICATION MODE (REMOTE CLIENT-SERVER MODE)

Replication mode is a clustered cache mode in Red Hat JBoss Data Grid. Replication mode can be added to any cache container using the following procedure.

#### Procedure 8.1. The *replicated-cache* Element

```
<cache-container name="clustered"
  default-cache="default"
  statistics="true">
  <transport executor="infinispan-transport"
    lock-timeout="60000"/>
  <replicated-cache name="default"
    mode="SYNC"
    start="EAGER"
    statistics="true">
    <transaction mode="NONE" />
  </replicated-cache>
</cache-container>
```



## IMPORTANT

JGroups must be appropriately configured for clustered mode before attempting to load this configuration.

The **replicated-cache** element configures settings for the distributed cache using the following parameters:

1. The **name** parameter provides a unique identifier for the cache.
2. The **mode** parameter sets the clustered cache mode. Valid values are **SYNC** (synchronous) and **ASYNC** (asynchronous).
3. The **start** parameter specifies whether the cache starts when the server starts up or when it is requested or deployed.
4. If **statistics** are enabled at the container level, per-cache statistics can be selectively disabled for caches that do not require monitoring by setting the **statistics** attribute to **false**.
5. The **transaction** element sets up the transaction mode for the replicated cache.



## IMPORTANT

In Remote Client-Server mode, the transactions element is set to **NONE** unless JBoss Data Grid is used in a compatibility mode where the cluster contains both JBoss Data Grid server and library instances. In this case, if transactions are configured in the library mode instance, they must also be configured in the server instance.

For details about the **cache-container** and **locking**, see the appropriate chapter.

[Report a bug](#)

## 8.4. CONFIGURE REPLICATION MODE (LIBRARY MODE)

The following procedure shows a replication mode configuration in Red Hat JBoss Data Grid's Library mode.

### Procedure 8.2. Replication Mode Configuration

#### 1. Configure the Clustering Element

```
<clustering mode="replication">
  <sync replTimeout="${TIME}" />
  <stateTransfer chunkSize="${SIZE}"
    fetchInMemoryState="{true/false}"
    awaitInitialTransfer="{true/false}"
    timeout="${TIME}" />
```

- a. The **clustering** element's **mode** parameter's value determines the clustering mode selected for the cache.

- b. The **sync** element's **replTimeout** parameter specifies the maximum time period in milliseconds for an acknowledgment after a remote call. If the time period ends without any acknowledgment, an exception is thrown.
- c. The **stateTransfer** element specifies how state is transferred when a node leaves or joins the cluster. It uses the following parameters:
  - i. The **chunkSize** parameter is the number of cache entries to be transferred in one chunk. The default **chunkSize** value is 512. The **chunkSize** depends on many parameters (for example, size of entries) and the best setting needs to be measured. To change the value for larger cache entries, smaller chunks are advisable and for smaller cache entries, increasing the **chunkSize** offers better performance.
  - ii. The **fetchInMemoryState** parameter when set to **true**, requests state information from neighboring caches on start up. This impacts the startup time for the cache. Set the default to **true** if there is no shared cache store behind to avoid cache inconsistency.
  - iii. The **awaitInitialTransfer** parameter causes the first call to method **CacheManager.getCache()** on the joiner node to block and wait until the joining is complete and the cache has finished receiving state from neighboring caches (if **fetchInMemoryState** is enabled). This option applies to distributed and replicated caches only and is enabled by default.
  - iv. The **timeout** parameter specifies the maximum time (in milliseconds) the cache waits for responses from neighboring caches with the requested states. If no response is received within the **timeout** period, the start up process aborts and an exception is thrown. As a result of the failed state transfer, the cache is not available on that instance.

## 2. Specify Transport Configuration

```
<global>
  <transport clusterName="{NAME}"
    distributedSyncTimeout="{TIME}"
    transportClass="{CLASS}" />
</global>
```

The **transport** element defines the transport configuration for the cache container as follows:

- a. The **clusterName** parameter specifies the name of the cluster. Nodes can only connect to clusters that share the same name.
- b. The **distributedSyncTimeout** parameter specifies the time to wait to acquire a lock on the distributed lock. This distributed lock ensures that a single cache can transfer state or rehash state at a time.
- c. The **transportClass** parameter specifies a class that represents a network transport for the cache container.

[Report a bug](#)

## 8.5. SYNCHRONOUS AND ASYNCHRONOUS REPLICATION

Replication mode can be synchronous or asynchronous depending on the problem being addressed.

- Synchronous replication blocks a thread or caller (for example on a **put()** operation) until the

modifications are replicated across all nodes in the cluster. By waiting for acknowledgments, synchronous replication ensures that all replications are successfully applied before the operation is concluded.

- Asynchronous replication operates significantly faster than synchronous replication because it does not need to wait for responses from nodes. Asynchronous replication performs the replication in the background and the call returns immediately. Errors that occur during asynchronous replication are written to a log. As a result, a transaction can be successfully completed despite the fact that replication of the transaction may not have succeeded on all the cache instances in the cluster.

[Report a bug](#)

### 8.5.1. Troubleshooting Asynchronous Replication Behavior

In some instances, a cache configured for asynchronous replication or distribution may wait for responses, which is synchronous behavior. This occurs because caches behave synchronously when both state transfers and asynchronous modes are configured. This synchronous behavior is a prerequisite for state transfer to operate as expected.

Use one of the following to remedy this problem:

- Disable state transfer and use a **ClusteredCacheLoader** to lazily look up remote state as and when needed.
- Enable state transfer and **REPL\_SYNC**. Use the Asynchronous API (for example, the `cache.putAsync(k, v)`) to activate 'fire-and-forget' capabilities.
- Enable state transfer and **REPL\_ASYNC**. All RPCs end up becoming synchronous, but client threads will not be held up if a replication queue is enabled (which is recommended for asynchronous mode).

[Report a bug](#)

## 8.6. THE REPLICATION QUEUE

In replication mode, Red Hat JBoss Data Grid uses a replication queue to replicate changes across nodes based on the following:

- Previously set intervals.
- The queue size exceeding the number of elements.
- A combination of previously set intervals and the queue size exceeding the number of elements.

The replication queue ensures that during replication, cache operations are transmitted in batches instead of individually. As a result, a lower number of replication messages are transmitted and fewer envelopes are used, resulting in improved JBoss Data Grid performance.

A disadvantage of using the replication queue is that the queue is periodically flushed based on the time or the queue size. Such flushing operations delay the realization of replication, distribution, or invalidation operations across cluster nodes. When the replication queue is disabled, the data is directly transmitted and therefore the data arrives at the cluster nodes faster.

A replication queue is used in conjunction with asynchronous mode.



[Report a bug](#)

### 8.6.1. Replication Queue Usage

When using the replication queue, do one of the following:

- Disable asynchronous marshalling.
- Set the *max-threads* count value to **1** for the **transport executor**. The **transport executor** is defined in **standalone.xml** or **clustered.xml** as follows:

```
<transport executor="infinispan-transport"/>
```

To implement either of these solutions, the replication queue must be in use in asynchronous mode. Asynchronous mode can be set, along with the queue timeout (*queue-flush-interval*, value is in milliseconds) and queue size (*queue-size*) as follows:

#### Example 8.1. Replication Queue in Asynchronous Mode

```
<replicated-cache name="asyncCache"
    start="EAGER"
    mode="ASYNC"
    batching="false"
    indexing="NONE"
    statistics="true"
    queue-size="1000"
    queue-flush-interval="500">
    <!-- Additional configuration information here -->
</replicated-cache>
```

The replication queue allows requests to return to the client faster, therefore using the replication queue together with asynchronous marshalling does not present any significant advantages.

[Report a bug](#)

## 8.7. ABOUT REPLICATION GUARANTEES

In a clustered cache, the user can receive synchronous replication guarantees as well as the parallelism associated with asynchronous replication. Red Hat JBoss Data Grid provides an asynchronous API for this purpose.

The asynchronous methods used in the API return Futures, which can be queried. The queries block the thread until a confirmation is received about the success of any network calls used.

[Report a bug](#)

## 8.8. REPLICATION TRAFFIC ON INTERNAL NETWORKS

Some cloud providers charge less for traffic over internal **IP** addresses than for traffic over public **IP** addresses, or do not charge at all for internal network traffic (for example, GoGrid). To take advantage of lower rates, you can configure Red Hat JBoss Data Grid to transfer replication traffic using the internal

network. With such a configuration, it is difficult to know the internal **IP** address you are assigned. JBoss Data Grid uses JGroups interfaces to solve this problem.

[Report a bug](#)

## CHAPTER 9. SET UP INVALIDATION MODE

### 9.1. ABOUT INVALIDATION MODE

Invalidation is a clustered mode that does not share any data, but instead removes potentially obsolete data from remote caches. Using this cache mode requires another, more permanent store for the data such as a database.

Red Hat JBoss Data Grid, in such a situation, is used as an optimization for a system that performs many read operations and prevents database usage each time a state is needed.

When invalidation mode is in use, data changes in a cache prompts other caches in the cluster to evict their outdated data from memory.

[Report a bug](#)

### 9.2. CONFIGURE INVALIDATION MODE (REMOTE CLIENT-SERVER MODE)

Invalidation mode is a clustered mode in Red Hat JBoss Data Grid. Invalidation mode can be added to any cache container using the following procedure:

#### Procedure 9.1. The `invalidation-cache` Element

```
<cache-container name="local"
  default-cache="default"
  statistics="true">
  <invalidation-cache name="default"
    mode="ASYNC"
    start="EAGER"
    statistics="true">
    <!-- Additional configuration information here -->
  </invalidation-cache>
</cache-container>
```

The `invalidation-cache` element configures settings for the distributed cache using the following parameters:

1. The **name** parameter provides a unique identifier for the cache.
2. The **mode** parameter sets the clustered cache mode. Valid values are **SYNC** (synchronous) and **ASYNC** (asynchronous).
3. The **start** parameter specifies whether the cache starts when the server starts up or when it is requested or deployed.
4. If **statistics** are enabled at the container level, per-cache statistics can be selectively disabled for caches that do not require monitoring by setting the **statistics** attribute to **false**.



## IMPORTANT

JGroups must be appropriately configured for clustered mode before attempting to load this configuration.

For details about the **cache-container**, **locking**, and **transaction** elements, see the appropriate chapter.

[Report a bug](#)

## 9.3. CONFIGURE INVALIDATION MODE (LIBRARY MODE)

The following procedure shows an invalidation mode cache configuration in Red Hat JBoss Data Grid's Library mode.

### Procedure 9.2. Invalidation Mode Configuration

#### 1. Configure the Clustering Element

```
<clustering mode="invalidation">
  <sync replTimeout="${TIME}" />
  <stateTransfer chunkSize="${SIZE}"
    fetchInMemoryState="{true/false}"
    awaitInitialTransfer="{true/false}"
    timeout="${TIME}" />
```

- a. The **clustering** element's **mode** parameter's value determines the clustering mode selected for the cache.
- b. The **sync** element's **replTimeout** parameter specifies the maximum time period in milliseconds for an acknowledgment after a remote call. If the time period ends without any acknowledgment, an exception is thrown.
- c. The **stateTransfer** element specifies how state is transferred when a node leaves or joins the cluster. It uses the following parameters:
  - i. The **chunkSize** parameter specifies the size of cache entry state batches to be transferred. If this value is greater than **0**, the value set is the size of chunks sent. If the value is less than **0**, all states are transferred at the same time.
  - ii. The **fetchInMemoryState** parameter when set to **true**, requests state information from neighboring caches on start up. This impacts the start up time for the cache.
  - iii. The **awaitInitialTransfer** parameter causes the first call to method **CacheManager.getCache()** on the joiner node to block and wait until the joining is complete and the cache has finished receiving state from neighboring caches (if **fetchInMemoryState** is enabled). This option applies to distributed and replicated caches only and is enabled by default.
  - iv. The **timeout** parameter specifies the maximum time (in milliseconds) the cache waits for responses from neighboring caches with the requested states. If no response is received within the **timeout** period, the start up process aborts and an exception is thrown.

## 2. Specify Transport Configuration

```
<global>
  <transport clusterName="${NAME}"
    distributedSyncTimeout="${TIME}"
    transportClass="${CLASS}" />
</global>
```

The **transport** element defines the transport configuration for the cache container as follows:

- a. The **clusterName** parameter specifies the name of the cluster. Nodes can only connect to clusters that share the same name.
- b. The **distributedSyncTimeout** parameter specifies the time to wait to acquire a lock on the distributed lock. This distributed lock ensures that a single cache can transfer state or rehash state at a time.
- c. The **transportClass** parameter specifies a class that represents a network transport for the cache container.

[Report a bug](#)

## 9.4. SYNCHRONOUS/ASYNCHRONOUS INVALIDATION

In Red Hat JBoss Data Grid's Library mode, invalidation operates either asynchronously or synchronously.

- Synchronous invalidation blocks the thread until all caches in the cluster have received invalidation messages and evicted the obsolete data.
- Asynchronous invalidation operates in a fire-and-forget mode that allows invalidation messages to be broadcast without blocking a thread to wait for responses.

[Report a bug](#)

## 9.5. THE L1 CACHE AND INVALIDATION

An invalidation message is generated each time a key is updated. This message is multicast to each node that contains data that corresponds to current L1 cache entries. The invalidation message ensures that each of these nodes marks the relevant entry as invalidated.

[Report a bug](#)

## PART V. REMOTE CLIENT-SERVER MODE INTERFACES

Red Hat JBoss Data Grid offers the following APIs to interact with the data grid in Remote Client-Server mode:

- The Asynchronous API (can only be used in conjunction with the Hot Rod Client in Remote Client-Server Mode)
- The REST Interface
- The Memcached Interface
- The Hot Rod Interface
  - The RemoteCache API

[Report a bug](#)

## CHAPTER 10. THE ASYNCHRONOUS API

In addition to synchronous API methods, Red Hat JBoss Data Grid also offers an asynchronous API that provides the same functionality in a non-blocking fashion.

The asynchronous method naming convention is similar to their synchronous counterparts, with **Async** appended to each method name. Asynchronous methods return a `Future` that contains the result of the operation.

For example, in a cache parameterized as `Cache(String, String)`, `Cache.put(String key, String value)` returns a `String`, while `Cache.putAsync(String key, String value)` returns a `Future(String)`.

[Report a bug](#)

### 10.1. ASYNCHRONOUS API BENEFITS

The asynchronous API does not block, which provides multiple benefits, such as:

- The guarantee of synchronous communication, with the added ability to handle failures and exceptions.
- Not being required to block a thread's operations until the call completes.

These benefits allow you to better harness the parallelism in your system, for example:

#### Example 10.1. Using the Asynchronous API

```
Set<Future<?>> futures = new HashSet<Future<?>>();
futures.add(cache.putAsync("key1", "value1"));
futures.add(cache.putAsync("key2", "value2"));
futures.add(cache.putAsync("key3", "value3"));
```

In the example, The following lines do not block the thread as they execute:

- `futures.add(cache.putAsync(key1, value1));`
- `futures.add(cache.putAsync(key2, value2));`
- `futures.add(cache.putAsync(key3, value3));`

The remote calls from the three `put` operations are executed in parallel. This is particularly useful when executed in distributed mode.

[Report a bug](#)

### 10.2. ABOUT ASYNCHRONOUS PROCESSES

For a typical write operation in Red Hat JBoss Data Grid, the following processes fall on the critical path, ordered from most resource-intensive to the least:

- Network calls


- Marshalling
- Writing to a cache store (optional)
- Locking

In JBoss Data Grid, using asynchronous methods removes network calls and marshalling from the critical path.

[Report a bug](#)

### 10.3. RETURN VALUES AND THE ASYNCHRONOUS API

When the asynchronous API is used in Red Hat JBoss Data Grid, the client code requires the asynchronous operation to return either the **Future** or the **NotifyingFuture** in order to query the previous value.



#### NOTE

**NotifyingFutures** are only available in JBoss Data Grid Library mode.

Call the following operation to obtain the result of an asynchronous operation. This operation blocks threads when called.

```
Future.get()
```

[Report a bug](#)



## CHAPTER 11. THE REST INTERFACE

Red Hat JBoss Data Grid provides a REST interface. The primary benefit of the REST API is that it allows for loose coupling between the client and server. The need for specific versions of client libraries and bindings is also eliminated. The REST API introduces an overhead, and requires a REST client or custom code to understand and create REST calls.

To interact with JBoss Data Grid's REST API only requires a HTTP client library. For Java, the Apache HTTP Commons Client is recommended. Alternatively, the `java.net` API can be used.

[Report a bug](#)

### 11.1. RUBY CLIENT CODE

The following code is an example of interacting with Red Hat JBoss Data Grid REST API using ruby. The provided code does not require any special libraries and standard `net/HTTP` libraries are sufficient.

#### Example 11.1. Using the REST API with Ruby

```
require 'net/http'

http = Net::HTTP.new('localhost', 8080)

#An example of how to create a new entry

http.post('/rest/MyData/MyKey', 'DATA_HERE', {"Content-Type" =>
"text/plain"})

#An example of using a GET operation to retrieve the key

puts http.get('/rest/MyData/MyKey').body

#An Example of using a PUT operation to overwrite the key

http.put('/rest/MyData/MyKey', 'MORE DATA', {"Content-Type" =>
"text/plain"})

#An example of Removing the remote copy of the key

http.delete('/rest/MyData/MyKey')

#An example of creating binary data

http.put('/rest/MyImages/Image.png',
File.read('/Users/michaelneale/logo.png'), {"Content-Type" =>
"image/png"})
```

[Report a bug](#)

### 11.2. USING JSON WITH RUBY EXAMPLE

#### Prerequisites

To use JavaScript Object Notation (JSON) with ruby to interact with Red Hat JBoss Data Grid's REST Interface, install the JSON Ruby library (see your platform's package manager or the Ruby documentation) and declare the requirement using the following code:

```
require 'json'
```

### Using JSON with Ruby

The following code is an example of how to use JavaScript Object Notation (JSON) in conjunction with Ruby to send specific data, in this case the name and age of an individual, using the **PUT** function.

```
data = {:name => "michael", :age => 42 }
http.put('/rest/Users/data/0', data.to_json, {"Content-Type" =>
"application/json"})
```

[Report a bug](#)

## 11.3. PYTHON CLIENT CODE

The following code is an example of interacting with the Red Hat JBoss Data Grid REST API using Python. The provided code requires only the standard HTTP library.

### Example 11.2. Using the REST API with Python

```
import httplib

#How to insert data

conn = httplib.HTTPConnection("localhost:8080")
data = "SOME DATA HERE \!" #could be string, or a file...
conn.request("POST", "/rest/Bucket/0", data, {"Content-Type":
"text/plain"})
response = conn.getresponse()
print response.status

#How to retrieve data

import httplib
conn = httplib.HTTPConnection("localhost:8080")
conn.request("GET", "/rest/Bucket/0")
response = conn.getresponse()
print response.status
print response.read()
```

[Report a bug](#)

## 11.4. JAVA CLIENT CODE

The following code is an example of interacting with Red Hat JBoss Data Grid REST API using Java.

### Example 11.3. Defining Imports

```
import java.io.BufferedReader;import java.io.IOException;
import java.io.InputStreamReader;import java.io.OutputStreamWriter;
import java.net.HttpURLConnection;import java.net.URL;
```

#### Example 11.4. Adding a String Value to a Cache

```
public class RestExample {

    /**
     * Method that puts a String value in cache.
     * @param urlServerAddress
     * @param value
     * @throws IOException
     */

    public void putMethod(String urlServerAddress, String value) throws
    IOException {
        System.out.println("-----");
        System.out.println("Executing PUT");
        System.out.println("-----");
        URL address = new URL(urlServerAddress);
        System.out.println("executing request " + urlServerAddress);
        HttpURLConnection connection = (HttpURLConnection)
address.openConnection();
        System.out.println("Executing put method of value: " + value);
        connection.setRequestMethod("PUT");
        connection.setRequestProperty("Content-Type", "text/plain");
        connection.setDoOutput(true);

        OutputStreamWriter outputStreamWriter = new
OutputStreamWriter(connection.getOutputStream());
        outputStreamWriter.write(value);

        connection.connect();
        outputStreamWriter.flush();

        System.out.println("-----");
        System.out.println(connection.getResponseCode() + " " +
connection.getResponseMessage());
        System.out.println("-----");

        connection.disconnect();
    }
}
```

The following code is an example of a method used that reads a value specified in a URL using Java to interact with the JBoss Data Grid REST Interface.

#### Example 11.5. Get a String Value from a Cache

```
/**
 * Method that gets an value by a key in url as param value.
 * @param urlServerAddress
```

```

    * @return String value
    * @throws IOException
    */
public String getMethod(String urlServerAddress) throws IOException {
    String line = new String();
    StringBuilder stringBuilder = new StringBuilder();

    System.out.println("-----");
    System.out.println("Executing GET");
    System.out.println("-----");

    URL address = new URL(urlServerAddress);
    System.out.println("executing request " + urlServerAddress);

    HttpURLConnection connection = (HttpURLConnection)
address.openConnection();
    connection.setRequestMethod("GET");
    connection.setRequestProperty("Content-Type", "text/plain");
    connection.setDoOutput(true);

    BufferedReader bufferedReader = new BufferedReader(new
InputStreamReader(connection.getInputStream()));

    connection.connect();

    while ((line = bufferedReader.readLine()) != null) {
        stringBuilder.append(line + '\n');
    }

    System.out.println("Executing get method of value: " +
stringBuilder.toString());

    System.out.println("-----");
    System.out.println(connection.getResponseCode() + " " +
connection.getResponseMessage());
    System.out.println("-----");

    connection.disconnect();

    return stringBuilder.toString();
}

```

### Example 11.6. Using a Java Main Method

```

/**
 * Main method example.
 * @param args
 * @throws IOException
 */
public static void main(String[] args) throws IOException {
//Note that the cache name is "cacheX"
RestExample restExample = new RestExample();
restExample.putMethod("http://localhost:8080/rest/cacheX/1", "Infinispan
REST Test");
}

```

```
restExample.getMethod("http://localhost:8080/rest/cacheX/1");
    }
}
```

[Report a bug](#)

## 11.5. THE REST INTERFACE CONNECTOR

Red Hat JBoss Data Grid supports three connector types, namely:

- The **hotrod-connector** element, which defines the configuration for a Hot Rod based connector.
- The **memcached-connector** element, which defines the configuration for a memcached based connector.
- The **rest-connector** element, which defines the configuration for a REST interface based connector.

The connector declaration enables a Hot Rod, Memcached, or REST server using a socket binding, which is declared within a **<socket-binding-group />**, and exposing the caches declared in the **local** container, using defaults for all other settings. The following examples show how to connect to Hot Rod, Memcached, and REST servers.

The REST connector differs from Hot Rod and Memcached because it requires a web subsystem. Therefore configurations such as socket-binding, worker threads, timeouts, etc, must be performed on the web subsystem. The following enables a REST server.

```
<rest-connector virtual-server="default-host"
  cache-container="local"
  security-domain="other"
  auth-method="BASIC"/>
```

See [Section 11.6, “Using the REST Interface”](#) for more information.

[Report a bug](#)

### 11.5.1. Configure REST Connectors

Use the following procedure to configure the **rest-connector** element in Red Hat JBoss Data Grid's Remote Client-Server mode.

#### Procedure 11.1. Configuring REST Connectors for Remote Client-Server Mode

```
<subsystem xmlns="urn:infinispan:server:endpoint:6.1">
  <rest-connector virtual-server="default-host"
    cache-container="local"
    context-path="{CONTEXT_PATH}"
    security-domain="{SECURITY_DOMAIN}"
    auth-method="{METHOD}"
    security-mode="{MODE}" />
</subsystem>
```

-

The **rest-connector** element specifies the configuration information for the REST connector.

1. The **virtual-server** parameter specifies the virtual server used by the REST connector. The default value for this parameter is **default-host**. This is an optional parameter.
2. The **cache-container** parameter names the cache container used by the REST connector. This is a mandatory parameter.
3. The **context-path** parameter specifies the context path for the REST connector. The default value for this parameter is an empty string (""). This is an optional parameter.
4. The **security-domain** parameter specifies that the specified domain, declared in the security subsystem, should be used to authenticate access to the REST endpoint. This is an optional parameter. If this parameter is omitted, no authentication is performed.
5. The **auth-method** parameter specifies the method used to retrieve credentials for the endpoint. The default value for this parameter is **BASIC**. Supported alternate values include **BASIC**, **DIGEST**, and **CLIENT-CERT**. This is an optional parameter.
6. The **security-mode** parameter specifies whether authentication is required only for write operations (such as PUT, POST and DELETE) or for read operations (such as GET and HEAD) as well. Valid values for this parameter are **WRITE** for authenticating write operations only, or **READ\_WRITE** to authenticate read and write operations. The default value for this parameter is **READ\_WRITE**.

[Report a bug](#)

## 11.6. USING THE REST INTERFACE

The REST Interface can be used in Red Hat JBoss Data Grid's Remote Client-Server mode to perform the following operations:

- Adding data
- Retrieving data
- Removing data

[Report a bug](#)

### 11.6.1. Adding Data Using REST

In Red Hat JBoss Data Grid's REST Interface, use the following methods to add data to the cache:

- HTTP **PUT** method
- HTTP **POST** method

When the **PUT** and **POST** methods are used, the body of the request contains this data, which includes any information added by the user.

Both the **PUT** and **POST** methods require a Content-Type header.

[Report a bug](#)

### 11.6.1.1. About PUT `/{cacheName}/{cacheKey}`

A **PUT** request from the provided URL form places the payload, from the request body in the targeted cache using the provided key. The targeted cache must exist on the server for this task to successfully complete.

As an example, in the following URL, the value **hr** is the cache name and **payRo11%2F3** is the key. The value **%2F** indicates that a `/` was used in the key.

```
http://someserver/rest/hr/payRo11%2F3
```

Any existing data is replaced and *Time-To-Live* and *Last-Modified* values are updated, if an update is required.



#### NOTE

A cache key that contains the value **%2F** to represent a `/` in the key (as in the provided example) can be successfully run if the server is started using the following argument:

```
-Dorg.apache.tomcat.util.buf.UDecoder.ALLOW_ENCODED_SLASH=true
```

[Report a bug](#)

### 11.6.1.2. About POST `/{cacheName}/{cacheKey}`

The **POST** method from the provided URL form places the payload (from the request body) in the targeted cache using the provided key. However, in a **POST** method, if a value in a cache/key exists, a **HTTP CONFLICT** status is returned and the content is not updated.

[Report a bug](#)

## 11.6.2. Retrieving Data Using REST

In Red Hat JBoss Data Grid's REST Interface, use the following methods to retrieve data from the cache:

- HTTP **GET** method.
- HTTP **HEAD** method.

[Report a bug](#)

### 11.6.2.1. About GET `/{cacheName}/{cacheKey}`

The **GET** method returns the data located in the supplied *cacheName*, matched to the relevant key, as the body of the response. The Content-Type header provides the type of the data. A browser can directly access the cache.

A unique entity tag (ETag) is returned for each entry along with a Last-Modified header which indicates the state of the data at the requested URL. ETags allow browsers (and other clients) to ask for data only in the case where it has changed (to save on bandwidth). ETag is a part of the HTTP standard and is supported by Red Hat JBoss Data Grid.

The type of content stored is the type returned. As an example, if a String was stored, a String is returned. An object which was stored in a serialized form must be manually deserialized.

[Report a bug](#)

### 11.6.2.2. About HEAD `/{cacheName}/{cacheKey}`

The **HEAD** method operates in a manner similar to the **GET** method, however returns no content (header fields are returned).

[Report a bug](#)

## 11.6.3. Removing Data Using REST

To remove data from Red Hat JBoss Data Grid using the REST interface, use the HTTP **DELETE** method to retrieve data from the cache. The **DELETE** method can:

- Remove a cache entry/value. (**DELETE** `/{cacheName}/{cacheKey}`)
- Remove all entries from a cache. (**DELETE** `/{cacheName}`)

[Report a bug](#)

### 11.6.3.1. About DELETE `/{cacheName}/{cacheKey}`

Used in this context (**DELETE** `/{cacheName}/{cacheKey}`), the **DELETE** method removes the key/value from the cache for the provided key.

[Report a bug](#)

### 11.6.3.2. About DELETE `/{cacheName}`

In this context (**DELETE** `/{cacheName}`), the **DELETE** method removes all entries in the named cache. After a successful **DELETE** operation, the HTTP status code **200** is returned.

[Report a bug](#)

### 11.6.3.3. Background Delete Operations

Set the value of the *performAsync* header to **true** to ensure an immediate return while the removal operation continues in the background.

[Report a bug](#)

## 11.6.4. REST Interface Operation Headers

The following table displays headers that are included in the Red Hat JBoss Data Grid REST Interface:

**Table 11.1. Header Types**



Headers	Mandatory/Optional	Values	Default Value	Details
Content-Type	Mandatory	-	-	If the Content-Type is set to <b>application/x-java-serialized-object</b> , it is stored as a Java object.
performAsync	Optional	True/False	-	If set to <b>true</b> , an immediate return occurs, followed by a replication of data to the cluster on its own. This feature is useful when dealing with bulk data inserts and large clusters.
timeToLiveSeconds	Optional	Numeric (positive and negative numbers)	<b>-1</b> (This value prevents expiration as a direct result of timeToLiveSeconds. Expiration values set elsewhere override this default value.)	Reflects the number of seconds before the entry in question is automatically deleted. Setting a negative value for timeToLiveSeconds provides the same result as the default value.
maxIdleTimeSeconds	Optional	Numeric (positive and negative numbers)	<b>-1</b> (This value prevents expiration as a direct result of maxIdleTimeSeconds. Expiration values set elsewhere override this default value.)	Contains the number of seconds after the last usage when the entry will be automatically deleted. Passing a negative value provides the same result as the default value.

The following combinations can be set for the *timeToLiveSeconds* and *maxIdleTimeSeconds* headers:

- If both the *timeToLiveSeconds* and *maxIdleTimeSeconds* headers are assigned the value **0**, the cache uses the default *timeToLiveSeconds* and *maxIdleTimeSeconds* values configured either using XML or programmatically.
- If only the *maxIdleTimeSeconds* header value is set to **0**, the *timeToLiveSeconds* value should be passed as the parameter (or the default **-1**, if the parameter is not present). Additionally, the *maxIdleTimeSeconds* parameter value defaults to the values configured either using XML or programmatically.
- If only the *timeToLiveSeconds* header value is set to **0**, expiration occurs immediately and the *maxIdleTimeSeconds* value is set to the value passed as a parameter (or the default **-1** if no parameter was supplied).

### Etag Based Headers

ETags (Entity Tags) are returned for each REST Interface entry, along with a *Last-Modified* header that indicates the state of the data at the supplied URL. ETags are used in HTTP operations to request data exclusively in cases where the data has changed to save bandwidth. The following headers support ETags (Entity Tags) based optimistic locking:

**Table 11.2. Entity Tag Related Headers**

Header	Algorithm	Example	Details
If-Match	If-Match = "If-Match" ":" ( "*"   1#entity-tag )	-	Used in conjunction with a list of associated entity tags to verify that a specified entity (that was previously obtained from a resource) remains current.
If-None-Match		-	Used in conjunction with a list of associated entity tags to verify that none of the specified entities (that was previously obtained from a resource) are current. This feature facilitates efficient updates of cached information when required and with minimal transaction overhead.

Header	Algorithm	Example	Details
If-Modified-Since	If-Modified-Since = "If-Modified-Since" ":" HTTP-date	If-Modified-Since: Sat, 29 Oct 1994 19:43:31 GMT	Compares the requested variant's last modification time and date with a supplied time and date value. If the requested variant has not been modified since the specified time and date, a <b>304</b> (not modified) response is returned without a message-body instead of an entity.
If-Unmodified-Since	If-Unmodified-Since = "If-Unmodified-Since" ":" HTTP-date	If-Unmodified-Since: Sat, 29 Oct 1994 19:43:31 GMT	Compares the requested variant's last modification time and date with a supplied time and date value. If the requested resources has not been modified since the supplied date and time, the specified operation is performed. If the requested resource has been modified since the supplied date and time, the operation is not performed and a <b>412</b> (Precondition Failed) response is returned.

[Report a bug](#)

## 11.7. REST INTERFACE SECURITY

### 11.7.1. Publish REST Endpoints as a Public Interface

Red Hat JBoss Data Grid's REST server operates as a management interface by default. To extend its operations to a public interface, alter the value of the *interface* parameter in the **socket-binding** element from **management** to **public** as follows:

```
<socket-binding name="http"
  interface="public"
  port="8080"/>
```

[Report a bug](#)

### 11.7.2. Enable Security for the REST Endpoint

Use the following procedure to enable security for the REST endpoint in Red Hat JBoss Data Grid.



## NOTE

The REST endpoint supports any of the JBoss Enterprise Application Platform security subsystem providers.

### Procedure 11.2. Enable Security for the REST Endpoint

To enable security for JBoss Data Grid when using the REST interface, make the following changes to `standalone.xml`:

#### 1. Specify Security Parameters

Ensure that the rest endpoint specifies a valid value for the *security-domain* and *auth-method* parameters. Recommended settings for these parameters are as follows:

```
<subsystem xmlns="urn:infinispan:server:endpoint:6.1">
    <rest-connector virtual-server="default-host"
        cache-container="local"
        security-domain="other"
        auth-method="BASIC"/>
</subsystem>
```

#### 2. Check Security Domain Declaration

Ensure that the security subsystem contains the corresponding security-domain declaration. For details about setting up security-domain declarations, see the JBoss Enterprise Application Platform 6 documentation.

#### 3. Add an Application User

Run the relevant script and enter the configuration settings to add an application user.

- a. Run the `adduser.sh` script (located in `$JDG_HOME/bin`).
  - On a Windows system, run the `adduser.bat` file (located in `$JDG_HOME/bin`) instead.
- b. When prompted about the type of user to add, select **Application User (application-users.properties)** by entering `b`.
- c. Accept the default value for realm (**ApplicationRealm**) by pressing the return key.
- d. Specify a username and password.
- e. When prompted for a role for the created user, enter **REST**.
- f. Ensure the username and application realm information is correct when prompted and enter "yes" to continue.

#### 4. Verify the Created Application User

Ensure that the created application user is correctly configured.

- a. Check the configuration listed in the `application-users.properties` file (located in `$JDG_HOME/standalone/configuration/`). The following is an example of what the correct configuration looks like in this file:

-

```
user1=2dc3eacfed8cf95a4a31159167b936fc
```

- b. Check the configuration listed in the **application-roles.properties** file (located in **\$JDG\_HOME/standalone/configuration/**). The following is an example of what the correct configuration looks like in this file:

```
user1=REST
```

## 5. Test the Server

Start the server and enter the following link in a browser window to access the REST endpoint:

```
http://localhost:8080/rest/namedCache
```



### NOTE

If testing using a GET request, a **405** response code is expected and indicates that the server was successfully authenticated.

[Report a bug](#)

## CHAPTER 12. THE MEMCACHED INTERFACE

Memcached is an in-memory caching system used to improve response and operation times for database-driven websites. The Memcached caching system defines a text based protocol called the Memcached protocol. The Memcached protocol uses in-memory objects or (as a last resort) passes to a persistent store such as a special memcached database.

Red Hat JBoss Data Grid offers a server that uses the Memcached protocol, removing the necessity to use Memcached separately with JBoss Data Grid. Additionally, due to JBoss Data Grid's clustering features, its data failover capabilities surpass those provided by Memcached.

[Report a bug](#)

### 12.1. ABOUT MEMCACHED SERVERS

Red Hat JBoss Data Grid contains a server module that implements the memcached protocol. This allows memcached clients to interact with one or multiple JBoss Data Grid based memcached servers.

The servers can be either:

- Standalone, where each server acts independently without communication with any other memcached servers.
- Clustered, where servers replicate and distribute data to other memcached servers.

[Report a bug](#)

### 12.2. MEMCACHED STATISTICS

The following table contains a list of valid statistics available using the memcached protocol in Red Hat JBoss Data Grid.

**Table 12.1. Memcached Statistics**

Statistic	Data Type	Details
uptime	32-bit unsigned integer.	Contains the time (in seconds) that the memcached instance has been available and running.
time	32-bit unsigned integer.	Contains the current time.
version	String	Contains the current version.
curr_items	32-bit unsigned integer.	Contains the number of items currently stored by the instance.
total_items	32-bit unsigned integer.	Contains the total number of items stored by the instance during its lifetime.

Statistic	Data Type	Details
cmd_get	64-bit unsigned integer	Contains the total number of get operation requests (requests to retrieve data).
cmd_set	64-bit unsigned integer	Contains the total number of set operation requests (requests to store data).
get_hits	64-bit unsigned integer	Contains the number of keys that are present from the keys requested.
get_misses	64-bit unsigned integer	Contains the number of keys that were not found from the keys requested.
delete_hits	64-bit unsigned integer	Contains the number of keys to be deleted that were located and successfully deleted.
delete_misses	64-bit unsigned integer	Contains the number of keys to be deleted that were not located and therefore could not be deleted.
incr_hits	64-bit unsigned integer	Contains the number of keys to be incremented that were located and successfully incremented
incr_misses	64-bit unsigned integer	Contains the number of keys to be incremented that were not located and therefore could not be incremented.
decr_hits	64-bit unsigned integer	Contains the number of keys to be decremented that were located and successfully decremented.
decr_misses	64-bit unsigned integer	Contains the number of keys to be decremented that were not located and therefore could not be decremented.
cas_hits	64-bit unsigned integer	Contains the number of keys to be compared and swapped that were found and successfully compared and swapped.

Statistic	Data Type	Details
cas_misses	64-bit unsigned integer	Contains the number of keys to be compared and swapped that were not found and therefore not compared and swapped.
cas_badval	64-bit unsigned integer	Contains the number of keys where a compare and swap occurred but the original value did not match the supplied value.
evictions	64-bit unsigned integer	Contains the number of eviction calls performed.
bytes_read	64-bit unsigned integer	Contains the total number of bytes read by the server from the network.
bytes_written	64-bit unsigned integer	Contains the total number of bytes written by the server to the network.

[Report a bug](#)

## 12.3. THE MEMCACHED INTERFACE CONNECTOR

Red Hat JBoss Data Grid supports three connector types, namely:

- The **hotrod-connector** element, which defines the configuration for a Hot Rod based connector.
- The **memcached-connector** element, which defines the configuration for a memcached based connector.
- The **rest-connector** element, which defines the configuration for a REST interface based connector.

The connector declaration enables a Hot Rod, Memcached, or REST server using a socket binding, which is declared within a `<socket-binding-group />`, and exposing the caches declared in the **local** container, using defaults for all other settings. The following examples show how to connect to Hot Rod, Memcached, and REST servers.

The following enables a Memcached server using the **memcached** socket binding, and exposes the **memcachedCache** cache declared in the **local** container, using defaults for all other settings.

```
<memcached-connector socket-binding="memcached"
    cache-container="local"/>
```

Due to the limitations in the Memcached protocol, only one cache can be exposed by a connector. To expose more than one cache, declare additional memcached-connectors on different socket-bindings. See [Section 12.3.1, “Configure Memcached Connectors”](#).



[Report a bug](#)

### 12.3.1. Configure Memcached Connectors

The following procedure describes the attributes used to configure the memcached connector within the **connectors** element in Red Hat JBoss Data Grid's Remote Client-Server Mode.

#### Procedure 12.1. Configuring the Memcached Connector in Remote Client-Server Mode

The **memcached-connector** element defines the configuration elements for use with memcached.

```
<subsystem xmlns="urn:infinispan:server:endpoint:6.1">
  <memcached-connector socket-binding="memcached"
    cache-container="local"
    worker-threads="{VALUE}"
    idle-timeout="{VALUE}"
    tcp-nodelay="{TRUE/FALSE}"
    send-buffer-size="{VALUE}"
    receive-buffer-size="{VALUE}" />
</subsystem>
```

1. The **socket-binding** parameter specifies the socket binding port used by the memcached connector. This is a mandatory parameter.
2. The **cache-container** parameter names the cache container used by the memcached connector. This is a mandatory parameter.
3. The **worker-threads** parameter specifies the number of worker threads available for the memcached connector. The default value for this parameter is 160. This is an optional parameter.
4. The **idle-timeout** parameter specifies the time (in milliseconds) the connector can remain idle before the connection times out. The default value for this parameter is **-1**, which means that no timeout period is set. This is an optional parameter.
5. The **tcp-nodelay** parameter specifies whether TCP packets will be delayed and sent out in batches. Valid values for this parameter are **true** and **false**. The default value for this parameter is **true**. This is an optional parameter.
6. The **send-buffer-size** parameter indicates the size of the send buffer for the memcached connector. The default value for this parameter is the size of the TCP stack buffer. This is an optional parameter.
7. The **receive-buffer-size** parameter indicates the size of the receive buffer for the memcached connector. The default value for this parameter is the size of the TCP stack buffer. This is an optional parameter.

[Report a bug](#)

## 12.4. MEMCACHED INTERFACE SECURITY

### 12.4.1. Publish Memcached Endpoints as a Public Interface

Red Hat JBoss Data Grid's memcached server operates as a management interface by default. To extend its operations to a public interface, alter the value of the ***interface*** parameter in the **socket-binding** element from **management** to **public** as follows:

```
<socket-binding name="memcached"  
  interface="public"  
  port="11211" />
```

[Report a bug](#)

## CHAPTER 13. THE HOT ROD INTERFACE

### 13.1. ABOUT HOT ROD

Hot Rod is a binary TCP client-server protocol used in Red Hat JBoss Data Grid. It was created to overcome deficiencies in other client/server protocols, such as Memcached.

Hot Rod will failover on a server cluster that undergoes a topology change. Hot Rod achieves this by providing regular updates to clients about the cluster topology.

Hot Rod enables clients to do smart routing of requests in partitioned or distributed JBoss Data Grid server clusters. To do this, Hot Rod allows clients to determine the partition that houses a key and then communicate directly with the server that has the key. This functionality relies on Hot Rod updating the cluster topology with clients, and that the clients use the same consistent hash algorithm as the servers.

JBoss Data Grid contains a server module that implements the Hot Rod protocol. The Hot Rod protocol facilitates faster client and server interactions in comparison to other text-based protocols and allows clients to make decisions about load balancing, failover and data location operations.

[Report a bug](#)

### 13.2. THE BENEFITS OF USING HOT ROD OVER MEMCACHED

Red Hat JBoss Data Grid offers a choice of protocols for allowing clients to interact with the server in a Remote Client-Server environment. When deciding between using memcached or Hot Rod, the following should be considered.

#### Memcached

The memcached protocol causes the server endpoint to use the **memcached text wire protocol**. The **memcached wire protocol** has the benefit of being commonly used, and is available for almost any platform. All of JBoss Data Grid's functions, including clustering, state sharing for scalability, and high availability, are available when using memcached.

However the memcached protocol lacks dynamicity, resulting in the need to manually update the list of server nodes on your clients in the event one of the nodes in a cluster fails. Also, memcached clients are not aware of the location of the data in the cluster. This means that they will request data from a non-owner node, incurring the penalty of an additional request from that node to the actual owner, before being able to return the data to the client. This is where the Hot Rod protocol is able to provide greater performance than memcached.

#### Hot Rod

JBoss Data Grid's Hot Rod protocol is a binary wire protocol that offers all the capabilities of memcached, while also providing better scaling, durability, and elasticity.

The Hot Rod protocol does not need the hostnames and ports of each node in the remote cache, whereas memcached requires these parameters to be specified. Hot Rod clients automatically detect changes in the topology of clustered Hot Rod servers; when new nodes join or leave the cluster, clients update their Hot Rod server topology view. Consequently, Hot Rod provides ease of configuration and maintenance, with the advantage of dynamic load balancing and failover.

Additionally, the Hot Rod wire protocol uses smart routing when connecting to a distributed cache. This involves sharing a consistent hash algorithm between the server nodes and clients, resulting in faster read and writing capabilities than memcached.



## WARNING

When using JCache over Hot Rod it is not possible to create remote clustered caches, as the operation is executed on a single node as opposed to the entire cluster; however, once a cache has been created on the cluster it may be obtained using the `cacheManager.getCache` method.

It is recommended to create caches using either configuration files, JON, or the CLI.

[Report a bug](#)

## 13.3. HOT ROD HASH FUNCTIONS

Hot Rod uses the same algorithm as on the server. The Hot Rod client always connects to the primary owner of the key, which is the first node in the list of owners. For more information about consistent hashing in Red Hat JBoss Data Grid, see [Section 7.2, “Distribution Mode’s Consistent Hash Algorithm”](#).

[Report a bug](#)

## 13.4. THE HOT ROD INTERFACE CONNECTOR

Red Hat JBoss Data Grid supports three connector types, namely:

- The **hotrod-connector** element, which defines the configuration for a Hot Rod based connector.
- The **memcached-connector** element, which defines the configuration for a memcached based connector.
- The **rest-connector** element, which defines the configuration for a REST interface based connector.

The connector declaration enables a Hot Rod, Memcached, or REST server using a socket binding, which is declared within a **<socket-binding-group />**, and exposing the caches declared in the **local** container, using defaults for all other settings. The following examples show how to connect to Hot Rod, Memcached, and REST servers.

The following enables a Hot Rod server using the **hotrod** socket binding.

```
<hotrod-connector socket-binding="hotrod"
  cache-container="local" />
```

The connector creates a supporting topology cache with default settings. These settings can be tuned by adding the **<topology-state-transfer />** child element to the connector as follows:


```
<hotrod-connector socket-binding="hotrod"
  cache-container="local">
  <topology-state-transfer lazy-retrieval="false"
```

```

        lock-timeout="1000"
        replication-timeout="5000" />
</hotrod-connector>

```

The Hot Rod connector can be tuned with additional settings. See [Section 13.4.1, “Configure Hot Rod Connectors”](#) for more information on how to configure the Hot Rod connector.



## NOTE

The Hot Rod connector can be secured using SSL. See the *Hot Rod Authentication Using SASL* section of the *Developer Guide* for more information.

[Report a bug](#)

### 13.4.1. Configure Hot Rod Connectors

The following procedure describes the attributes used to configure the Hot Rod connector in Red Hat JBoss Data Grid's Remote Client-Server Mode. Both the **hotrod-connector** and **topology-state-transfer** elements must be configured based on the following procedure.

#### Procedure 13.1. Configuring Hot Rod Connectors for Remote Client-Server Mode

```

<subsystem xmlns="urn:infinispan:server:endpoint:6.1">
  <hotrod-connector socket-binding="hotrod"
    cache-container="local"
    worker-threads="${VALUE}"
    idle-timeout="${VALUE}"
    tcp-nodelay="${TRUE/FALSE}"
    send-buffer-size="${VALUE}"
    receive-buffer-size="${VALUE}" />
  <topology-state-transfer lock-timeout="${MILLISECONDS}"
    replication-timeout="${MILLISECONDS}"
    external-host="${HOSTNAME}"
    external-port="${PORT}"
    lazy-retrieval="${TRUE/FALSE}"
    await-initial-transfer="${TRUE/FALSE}" />
</subsystem>

```

1. The **hotrod-connector** element defines the configuration elements for use with Hot Rod.
  - a. The **socket-binding** parameter specifies the socket binding port used by the Hot Rod connector. This is a mandatory parameter.
  - b. The **cache-container** parameter names the cache container used by the Hot Rod connector. This is a mandatory parameter.
  - c. The **worker-threads** parameter specifies the number of worker threads available for the Hot Rod connector. The default value for this parameter is **160**. This is an optional parameter.
  - d. The **idle-timeout** parameter specifies the time (in milliseconds) the connector can remain idle before the connection times out. The default value for this parameter is **-1**, which means that no timeout period is set. This is an optional parameter.

- e. The ***tcp-nodelay*** parameter specifies whether TCP packets will be delayed and sent out in batches. Valid values for this parameter are **true** and **false**. The default value for this parameter is **true**. This is an optional parameter.
  - f. The ***send-buffer-size*** parameter indicates the size of the send buffer for the Hot Rod connector. The default value for this parameter is the size of the TCP stack buffer. This is an optional parameter.
  - g. The ***receive-buffer-size*** parameter indicates the size of the receive buffer for the Hot Rod connector. The default value for this parameter is the size of the TCP stack buffer. This is an optional parameter.
2. The **topology-state-transfer** element specifies the topology state transfer configurations for the Hot Rod connector. This element can only occur once within a **hotrod-connector** element.
- a. The ***lock-timeout*** parameter specifies the time (in milliseconds) after which the operation attempting to obtain a lock times out. The default value for this parameter is **10** seconds. This is an optional parameter.
  - b. The ***replication-timeout*** parameter specifies the time (in milliseconds) after which the replication operation times out. The default value for this parameter is **10** seconds. This is an optional parameter.
  - c. The ***external-host*** parameter specifies the hostname sent by the Hot Rod server to clients listed in the topology information. The default value for this parameter is the host address. This is an optional parameter.
  - d. The ***external-port*** parameter specifies the port sent by the Hot Rod server to clients listed in the topology information. The default value for this parameter is the configured port. This is an optional parameter.
  - e. The ***lazy-retrieval*** parameter indicates whether the Hot Rod connector will carry out retrieval operations lazily. The default value for this parameter is **true**. This is an optional parameter.
  - f. The ***await-initial-transfer*** parameter specifies whether the initial state retrieval happens immediately at startup. This parameter only applies when ***lazy-retrieval*** is set to **false**. This default value for this parameter is **true**.

[Report a bug](#)

## 13.5. HOT ROD HEADERS

### 13.5.1. Hot Rod Header Data Types

All keys and values used for Hot Rod in Red Hat JBoss Data Grid are stored as byte arrays. Certain header values, such as those for REST and Memcached, are stored using the following data types instead:

**Table 13.1. Header Data Types**

Data Type	Size	Details
vInt	Between 1-5 bytes.	Unsigned variable length integer values.
vLong	Between 1-9 bytes.	Unsigned variable length long values.
string	-	Strings are always represented using UTF-8 encoding.

[Report a bug](#)

### 13.5.2. Request Header

When using Hot Rod to access Red Hat JBoss Data Grid, the contents of the request header consist of the following:

**Table 13.2. Request Header Fields**

Field Name	Data Type/Size	Details
Magic	1 byte	Indicates whether the header is a request header or response header.
Message ID	vLong	Contains the message ID. Responses use this unique ID when responding to a request. This allows Hot Rod clients to implement the protocol in an asynchronous manner.
Version	1 byte	Contains the Hot Rod server version.
Opcode	1 byte	Contains the relevant operation code. In a request header, opcode can only contain the request operation codes.
Cache Name Length	vInt	Stores the length of the cache name. If Cache Name Length is set to <b>0</b> and no value is supplied for Cache Name, the operation interacts with the default cache.

Field Name	Data Type/Size	Details
Cache Name	string	Stores the name of the target cache for the specified operation. This name must match the name of a predefined cache in the cache configuration file.
Flags	vInt	Contains a numeric value of variable length that represents flags passed to the system. Each bit represents a flag, except the most significant bit, which is used to determine whether more bytes must be read. Using a bit to represent each flag facilitates the representation of flag combinations in a condensed manner.
Client Intelligence	1 byte	Contains a value that indicates the client capabilities to the server.
Topology ID	vInt	Contains the last known view ID in the client. Basic clients supply the value 0 for this field. Clients that support topology or hash information supply the value 0 until the server responds with the current view ID, which is subsequently used until a new view ID is returned by the server to replace the current view ID.
Transaction Type	1 byte	Contains a value that represents one of two known transaction types. Currently, the only supported value is 0.
Transaction ID	byte-array	Contains a byte array that uniquely identifies the transaction associated with the call. The transaction type determines the length of this byte array. If the value for <b>Transaction Type</b> was set to 0, no Transaction ID is present.

[Report a bug](#)

### 13.5.3. Response Header



When using Hot Rod to access Red Hat JBoss Data Grid, the contents of the response header consist of the following:

**Table 13.3. Response Header Fields**

Field Name	Data Type	Details
Magic	1 byte	Indicates whether the header is a request or response header.
Message ID	vLong	Contains the message ID. This unique ID is used to pair the response with the original request. This allows Hot Rod clients to implement the protocol in an asynchronous manner.
Opcode	1 byte	Contains the relevant operation code. In a response header, opcode can only contain the response operation codes.
Status	1 byte	Contains a code that represents the status of the response.
Topology Change Marker	1 byte	Contains a marker byte that indicates whether the response is included in the topology change information.

[Report a bug](#)

### 13.5.4. Topology Change Headers

When using Hot Rod to access Red Hat JBoss Data Grid, response headers respond to changes in the cluster or view formation by looking for clients that can distinguish between different topologies or hash distributions. The Hot Rod server compares the current *topology ID* and the *topology ID* sent by the client and, if the two differ, it returns a new *topology ID*.

[Report a bug](#)

#### 13.5.4.1. Topology Change Marker Values

The following is a list of valid values for the *Topology Change Marker* field in a response header:

**Table 13.4. Topology Change Marker Field Values**

Value	Details
0	No topology change information is added.

Value	Details
1	Topology change information is added.

[Report a bug](#)

### 13.5.4.2. Topology Change Headers for Topology-Aware Clients

The response header sent to topology-aware clients when a topology change is returned by the server includes the following elements:

**Table 13.5. Topology Change Header Fields**

Response Header Fields	Data Type/Size	Details
Response Header with Topology Change Marker	-	-
Topology ID	vInt	-
Num Servers in Topology	vInt	Contains the number of Hot Rod servers running in the cluster. This value can be a subset of the entire cluster if only some nodes are running Hot Rod servers.
mX: Host/IP Length	vInt	Contains the length of the hostname or IP address of an individual cluster member. Variable length allows this element to include hostnames, IPv4 and IPv6 addresses.
mX: Host/IP Address	string	Contains the hostname or IP address of an individual cluster member. The Hot Rod client uses this information to access the individual cluster member.
mX: Port	Unsigned Short. 2 bytes	Contains the port used by Hot Rod clients to communicate with the cluster member.

The three entries with the prefix **mX**, are repeated for each server in the topology. The first server in the topology's information fields will be prefixed with **m1** and the numerical value is incremented by one for each additional server till the value of **X** equals the number of servers specified in the *num servers in topology* field.

[Report a bug](#)

### 13.5.4.3. Topology Change Headers for Hash Distribution-Aware Clients

The response header sent to clients when a topology change is returned by the server includes the following elements:

**Table 13.6. Topology Change Header Fields**

Field	Data Type/Size	Details
Response Header with Topology Change Marker	-	-
Topology ID	vInt	-
Number Key Owners	Unsigned short. 2 bytes.	Contains the number of globally configured copies for each distributed key. Contains the value 0 if distribution is not configured on the cache.
Hash Function Version	1 byte	Contains a pointer to the hash function in use. Contains the value 0 if distribution is not configured on the cache.
Hash Space Size	vInt	Contains the modulus used by JBoss Data Grid for all module arithmetic related to hash code generation. Clients use this information to apply the correct hash calculations to the keys. Contains the value 0 if distribution is not configured on the cache.
Number servers in topology	vInt	Contains the number of <b>Hot Rod</b> servers running in the cluster. This value can be a subset of the entire cluster if only some nodes are running <b>Hot Rod</b> servers. This value also represents the number of host to port pairings included in the header.
Number Virtual Nodes Owners	vInt	Contains the number of configured virtual nodes. Contains the value 0 if no virtual nodes are configured or if distribution is not configured on the cache.

Field	Data Type/Size	Details
mX: Host/IP Length	vInt	Contains the length of the hostname or <b>IP</b> address of an individual cluster member. Variable length allows this element to include hostnames, <b>IPv4</b> and <b>IPv6</b> addresses.
mX: Host/IP Address	string	Contains the hostname or <b>IP</b> address of an individual cluster member. The <b>Hot Rod</b> client uses this information to access the individual cluster member.
mX: Port	Unsigned short. 2 bytes.	Contains the port used by <b>Hot Rod</b> clients to communicate with the cluster member.
mX: Hashcode	4 bytes.	

The three entries with the prefix **mX**, are repeated for each server in the topology. The first server in the topology's information fields will be prefixed with **m1** and the numerical value is incremented by one for each additional server till the value of **X** equals the number of servers specified in the ***num servers in topology*** field.

[Report a bug](#)

## 13.6. HOT ROD OPERATIONS

The following are valid operations when using Hot Rod protocol 1.3 to interact with Red Hat JBoss Data Grid:

- BulkGetKeys
- BulkGet
- Clear
- ContainsKey
- Get
- GetWithMetadata
- Ping
- PutIfAbsent
- Put
- Query

- `RemoveIfUnmodified`
- `Remove`
- `ReplaceIfUnmodified`
- `Replace`
- `Stats`



### IMPORTANT

When using the `RemoteCache` API to call the Hot Rod client's `Put`, `PutIfAbsent`, `Replace`, and `ReplaceWithVersion` operations, if `lifespan` is set to a value greater than 30 days, the value is treated as UNIX time and represents the number of seconds since the date 1/1/1970.

[Report a bug](#)

## 13.6.1. Hot Rod BulkGet Operation

A Hot Rod `BulkGet` operation uses the following request format:

**Table 13.7. BulkGet Operation Request Format**

Field	Data Type	Details
Header	-	-
Entry Count	<code>vInt</code>	Contains the maximum number of Red Hat JBoss Data Grid entries to be returned by the server. The entry is the key and value pair.

The response header for this operation contains one of the following response statuses:

**Table 13.8. BulkGet Operation Response Format**

Field	Data Type	Details
Header	-	-
More	<code>vInt</code>	Represents if more entries must be read from the stream. While <b>More</b> is set to <b>1</b> , additional entries follow until the value of <b>More</b> is set to <b>0</b> , which indicates the end of the stream.
Key Size	-	Contains the size of the key.

Field	Data Type	Details
Key	-	Contains the key value.
Value Size	-	Contains the size of the value.
Value	-	Contains the value.

For each entry that was requested, a **More**, **Key Size**, **Key**, **Value Size** and **Value** entry is appended to the response.

[Report a bug](#)

### 13.6.2. Hot Rod BulkGetKeys Operation

A Hot Rod **BulkGetKeys** operation uses the following request format:

**Table 13.9. BulkGetKeys Operation Request Format**

Field	Data Type	Details
Header	variable	Request header.

Field	Data Type	Details
Scope	vInt	<ul style="list-style-type: none"> <li>• <b>0</b> = Default Scope - This scope is used by <b>RemoteCache.keySet()</b> method. If the remote cache is a distributed cache, the server launches a map/reduce operation to retrieve all keys from all of the nodes (A topology-aware Hot Rod Client could be load balancing the request to any one node in the cluster). Otherwise, it will get keys from the cache instance local to the server receiving the request, as the keys must be the same across all nodes in a replicated cache.</li> <li>• <b>1</b> = Global Scope - This scope behaves the same to Default Scope.</li> <li>• <b>2</b> = Local Scope - In situations where the remote cache is a distributed cache, the server will not launch a map/reduce operation to retrieve keys from all nodes. Instead, it will only get keys local from the cache instance local to the server receiving the request.</li> </ul>

The response header for this operation contains one of the following response statuses:

**Table 13.10. BulkGetKeys Operation Response Format**

Field	Data Type	Details
Header	variable	Response header
Response status	1 byte	<b>0x00</b> = success, data follows.

Field	Data Type	Details
More	1 byte	One byte representing whether more keys need to be read from the stream. When set to <b>1</b> an entry follows, when set to <b>0</b> , it is the end of stream and no more entries are left to read.
Key 1 Length	vInt	Length of key
Key 1	Byte array	Retrieved key.
More	1 byte	-
Key 2 Length	vInt	-
Key 2	byte array	-
...etc		

[Report a bug](#)

### 13.6.3. Hot Rod Clear Operation

The `clear` operation format includes only a header.

Valid response statuses for this operation are as follows:

**Table 13.11. Clear Operation Response**

Response Status	Details
0x00	Red Hat JBoss Data Grid was successfully cleared.

[Report a bug](#)

### 13.6.4. Hot Rod ContainsKey Operation

A Hot Rod `containsKey` operation uses the following request format:

**Table 13.12. ContainsKey Operation Request Format**

Field	Data Type	Details
Header	-	-



Field	Data Type	Details
Key Length	vInt	Contains the length of the key. The <b>vInt</b> data type is used because of its size (up to <b>5</b> bytes), which is larger than the size of <b>Integer.MAX_VALUE</b> . However, Java disallows single array sizes to exceed the size of <b>Integer.MAX_VALUE</b> . As a result, this <b>vInt</b> is also limited to the maximum size of <b>Integer.MAX_VALUE</b> .
Key	Byte array	Contains a key, the corresponding value of which is requested.

The response header for this operation contains one of the following response statuses:

**Table 13.13. ContainsKey Operation Response Format**

Response Status	Details
0x00	Successful operation.
0x02	The key does not exist.

The response for this operation is empty.

[Report a bug](#)

### 13.6.5. Hot Rod Get Operation

A Hot Rod **Get** operation uses the following request format:

**Table 13.14. Get Operation Request Format**

Field	Data Type	Details
Header	-	-

Field	Data Type	Details
Key Length	vInt	Contains the length of the key. The <b>vInt</b> data type is used because of its size (up to <b>5</b> bytes), which is larger than the size of <b>Integer.MAX_VALUE</b> . However, Java disallows single array sizes to exceed the size of <b>Integer.MAX_VALUE</b> . As a result, this vInt is also limited to the maximum size of <b>Integer.MAX_VALUE</b> .
Key	Byte array	Contains a key, the corresponding value of which is requested.

The response header for this operation contains one of the following response statuses:

**Table 13.15. Get Operation Response Format**

Response Status	Details
0x00	Successful operation.
0x02	The key does not exist.

The format of the **get** operation's response when the key is found is as follows:

**Table 13.16. Get Operation Response Format**

Field	Data Type	Details
Header	-	-
Value Length	vInt	Contains the length of the value.
Value	Byte array	Contains the requested value.

[Report a bug](#)

### 13.6.6. Hot Rod GetWithMetadata Operation

A Hot Rod **GetWithMetadata** operation uses the following request format:

**Table 13.17. GetWithMetadata Operation Request Format**

Field	Data Type	Details
Header	variable	Request header.
Key Length	vInt	Length of key. Note that the size of a vInt can be up to five bytes, which theoretically can produce bigger numbers than <b>Integer.MAX_VALUE</b> . However, Java cannot create a single array that is bigger than <b>Integer.MAX_VALUE</b> , hence the protocol limits vInt array lengths to <b>Integer.MAX_VALUE</b> .
Key	byte array	Byte array containing the key whose value is being requested.

The response header for this operation contains one of the following response statuses:

**Table 13.18. GetWithMetadata Operation Response Format**

Field	Data Type	Details
Header	variable	Response header.
Response status	1 byte	<b>0x00</b> = success, if key retrieved. <b>0x02</b> = if key does not exist.
Flag	1 byte	A flag indicating whether the response contains expiration information. The value of the flag is obtained as a bitwise OR operation between <b>INFINITE_LIFESPAN (0x01)</b> and <b>INFINITE_MAXIDLE (0x02)</b> .
Created	Long	(optional) a Long representing the timestamp when the entry was created on the server. This value is returned only if the flag's <b>INFINITE_LIFESPAN</b> bit is not set.

Field	Data Type	Details
Lifespan	vInt	(optional) a vInt representing the lifespan of the entry in seconds. This value is returned only if the flag's <b>INFINITE_LIFESPAN</b> bit is not set.
LastUsed	Long	(optional) a Long representing the timestamp when the entry was last accessed on the server. This value is returned only if the flag's <b>INFINITE_MAXIDLE</b> bit is not set.
MaxIdle	vInt	(optional) a vInt representing the maxIdle of the entry in seconds. This value is returned only if the flag's <b>INFINITE_MAXIDLE</b> bit is not set.
Entry Version	8 bytes	Unique value of an existing entry modification. The protocol does not mandate that entry_version values are sequential, however they need to be unique per update at the key level.
Value Length	vInt	If success, length of value.
Value	byte array	If success, the requested value.

[Report a bug](#)

### 13.6.7. Hot Rod Ping Operation

The **ping** is an application level request to check for server availability.

Valid response statuses for this operation are as follows:

**Table 13.19. Ping Operation Response**

Response Status	Details
0x00	Successful ping without any errors.

[Report a bug](#)

### 13.6.8. Hot Rod Put Operation

The **put** operation request format includes the following:

**Table 13.20.**


Field	Data Type	Details
Header	-	-
Key Length	-	Contains the length of the key.
Key	Byte array	Contains the key value.
Lifespan	vInt	Contains the number of seconds before the entry expires. If the number of seconds exceeds thirty days, the value is treated as UNIX time (i.e. the number of seconds since the date <b>1/1/1970</b> ) as the entry lifespan. When set to the value <b>0</b> , the entry will never expire.
Max Idle	vInt	Contains the number of seconds an entry is allowed to remain idle before it is evicted from the cache. If this entry is set to <b>0</b> , the entry is allowed to remain idle indefinitely without being evicted due to the <b>max idle</b> value.
Value Length	vInt	Contains the length of the value.
Value	Byte array	The requested value.

The following are the value response values returned from this operation:

**Table 13.21.**

Response Status	Details
0x00	The value was successfully stored.

An empty response is the default response for this operation. However, if **ForceReturnPreviousValue** is passed, the previous value and key are returned. If the previous key and value do not exist, the value length would contain the value **0**.



## IMPORTANT

When using the RemoteCache API to call the Hot Rod client's **Put**, **PutIfAbsent**, **Replace**, and **ReplaceWithVersion** operations, if lifespan is set to a value greater than 30 days, the value is treated as UNIX time and represents the number of seconds since the date 1/1/1970.

[Report a bug](#)

### 13.6.9. Hot Rod PutIfAbsent Operation

The `putIfAbsent` operation request format includes the following:

**Table 13.22. PutIfAbsent Operation Request Fields**

Field	Data Type	Details
Header	-	-
Key Length	vInt	Contains the length of the key.
Key	Byte array	Contains the key value.
Lifespan	vInt	Contains the number of seconds before the entry expires. If the number of seconds exceeds thirty days, the value is treated as UNIX time (i.e. the number of seconds since the date <b>1/1/1970</b> ) as the entry lifespan. When set to the value <b>0</b> , the entry will never expire.
Max Idle	vInt	Contains the number of seconds an entry is allowed to remain idle before it is evicted from the cache. If this entry is set to <b>0</b> , the entry is allowed to remain idle indefinitely without being evicted due to the max idle value.
Value Length	vInt	Contains the length of the value.
Value	Byte array	Contains the requested value.

The following are the value response values returned from this operation:

**Table 13.23.**

Response Status	Details
0x00	The value was successfully stored.
0x01	The key was present, therefore the value was not stored. The current value of the key is returned.

An empty response is the default response for this operation. However, if **ForceReturnPreviousValue** is passed, the previous value and key are returned. If the previous key and value do not exist, the value length would contain the value **0**.



### IMPORTANT

When using the RemoteCache API to call the Hot Rod client's **Put**, **PutIfAbsent**, **Replace**, and **ReplaceWithVersion** operations, if lifespan is set to a value greater than 30 days, the value is treated as UNIX time and represents the number of seconds since the date 1/1/1970.

[Report a bug](#)

## 13.6.10. Hot Rod Query Operation

The **Query** operation request format includes the following:

**Table 13.24. Query Operation Request Fields**

Field	Data Type	Details
Header	variable	Request header.
Query Length	vInt	The length of the Protobuf encoded query object.
Query	Byte array	Byte array containing the Protobuf encoded query object, having a length specified by previous field.

The following are the value response values returned from this operation:

**Table 13.25. Query Operation Response**

Response Status	Data	Details
Header	variable	Response header.
Response payload Length	vInt	The length of the Protobuf encoded response object.

Response Status	Data	Details
Response payload	Byte array	Byte array containing the Protobuf encoded response object, having a length specified by previous field.

[Report a bug](#)

### 13.6.11. Hot Rod Remove Operation

A **Hot Rod Remove** operation uses the following request format:

**Table 13.26. Remove Operation Request Format**

Field	Data Type	Details
Header	-	-
Key Length	vInt	Contains the length of the key. The <b>vInt</b> data type is used because of its size (up to <b>5</b> bytes), which is larger than the size of <b>Integer.MAX_VALUE</b> . However, Java disallows single array sizes to exceed the size of <b>Integer.MAX_VALUE</b> . As a result, this <b>vInt</b> is also limited to the maximum size of <b>Integer.MAX_VALUE</b> .
Key	Byte array	Contains a key, the corresponding value of which is requested.

The response header for this operation contains one of the following response statuses:

**Table 13.27. Remove Operation Response Format**

Response Status	Details
0x00	Successful operation.
0x02	The key does not exist.

Normally, the response header for this operation is empty. However, if **ForceReturnPreviousValue** is passed, the response header contains either:

- The value and length of the previous key.
- The value length **0** and the response status **0x02** to indicate that the key does not exist.



The remove operation's response header contains the previous value and the length of the previous value for the provided key if *ForceReturnPreviousValue* is passed. If the key does not exist or the previous value was null, the value length is 0.

[Report a bug](#)

### 13.6.12. Hot Rod RemoveIfUnmodified Operation

The **RemoveIfUnmodified** operation request format includes the following:

**Table 13.28. RemoveIfUnmodified Operation Request Fields**

Field	Data Type	Details
Header	-	-
Key Length	vInt	Contains the length of the key.
Key	Byte array	Contains the key value.
Entry Version	8 bytes	The version number for the entry.

The following are the value response values returned from this operation:

**Table 13.29. RemoveIfUnmodified Operation Response**

Response Status	Details
0x00	Returned status if the entry was replaced or removed.
0x01	Returns status if the entry replace or remove was unsuccessful because the key was modified.
0x02	Returns status if the key does not exist.

An empty response is the default response for this operation. However, if *ForceReturnPreviousValue* is passed, the previous value and key are returned. If the previous key and value do not exist, the value length would contain the value 0.

[Report a bug](#)

### 13.6.13. Hot Rod Replace Operation

The **replace** operation request format includes the following:

**Table 13.30. Replace Operation Request Fields**

Field	Data Type	Details
Header	-	-
Key Length	vInt	Contains the length of the key.
Key	Byte array	Contains the key value.
Lifespan	vInt	Contains the number of seconds before the entry expires. If the number of seconds exceeds thirty days, the value is treated as UNIX time (i.e. the number of seconds since the date <b>1/1/1970</b> ) as the entry lifespan. When set to the value <b>0</b> , the entry will never expire.
Max Idle	vInt	Contains the number of seconds an entry is allowed to remain idle before it is evicted from the cache. If this entry is set to <b>0</b> , the entry is allowed to remain idle indefinitely without being evicted due to the max idle value.
Value Length	vInt	Contains the length of the value.
Value	Byte array	Contains the requested value.

The following are the value response values returned from this operation:

**Table 13.31. Replace Operation Response**

Response Status	Details
0x00	The value was successfully stored.
0x01	The value was not stored because the key does not exist.

An empty response is the default response for this operation. However, if **ForceReturnPreviousValue** is passed, the previous value and key are returned. If the previous key and value do not exist, the value length would contain the value **0**.

**IMPORTANT**

When using the RemoteCache API to call the Hot Rod client's **Put**, **PutIfAbsent**, **Replace**, and **ReplaceWithVersion** operations, if lifespan is set to a value greater than 30 days, the value is treated as UNIX time and represents the number of seconds since the date 1/1/1970.

[Report a bug](#)

**13.6.14. Hot Rod ReplaceWithVersion Operation**

The **ReplaceWithVersion** operation request format includes the following:

**NOTE**

In the RemoteCache API, the Hot Rod **ReplaceWithVersion** operation uses the **ReplaceIfUnmodified** operation. As a result, these two operations are exactly the same in JBoss Data Grid.

**Table 13.32. ReplaceWithVersion Operation Request Fields**

Field	Data Type	Details
Header	-	-
Key Length	vInt	Contains the length of the key.
Key	Byte array	Contains the key value.
Lifespan	vInt	Contains the number of seconds before the entry expires. If the number of seconds exceeds thirty days, the value is treated as UNIX time (i.e. the number of seconds since the date <b>1/1/1970</b> ) as the entry lifespan. When set to the value <b>0</b> , the entry will never expire.
Max Idle	vInt	Contains the number of seconds an entry is allowed to remain idle before it is evicted from the cache. If this entry is set to <b>0</b> , the entry is allowed to remain idle indefinitely without being evicted due to the max idle value.
Entry Version	8 bytes	The version number for the entry.
Value Length	vInt	Contains the length of the value.

Field	Data Type	Details
Value	Byte array	Contains the requested value.

The following are the value response values returned from this operation:

**Table 13.33. ReplaceWithVersion Operation Response**

Response Status	Details
0x00	Returned status if the entry was replaced or removed.
0x01	Returns status if the entry replace or remove was unsuccessful because the key was modified.
0x02	Returns status if the key does not exist.

An empty response is the default response for this operation. However, if ***ForceReturnPreviousValue*** is passed, the previous value and key are returned. If the previous key and value do not exist, the value length would contain the value **0**.

[Report a bug](#)

### 13.6.15. Hot Rod Stats Operation

This operation returns a summary of all available statistics. For each returned statistic, a name and value is returned in both string and UTF-8 formats.

The following are supported statistics for this operation:

**Table 13.34. Stats Operation Request Fields**

Name	Details
timeSinceStart	Contains the number of seconds since Hot Rod started.
currentNumberOfEntries	Contains the number of entries that currently exist in the Hot Rod server.
totalNumberOfEntries	Contains the total number of entries stored in the Hot Rod server.
stores	Contains the number of put operations attempted.
retrievals	Contains the number of get operations attempted.
hits	Contains the number of get hits.

Name	Details
misses	Contains the number of get misses.
removeHits	Contains the number of remove hits.
removeMisses	Contains the number of removal misses.

The response header for this operation contains the following:

**Table 13.35. Stats Operation Response**

Name	Data Type	Details
Header	-	-
Number of Stats	vInt	Contains the number of individual statistics returned.
Name Length	vInt	Contains the length of the named statistic.
Name	string	Contains the name of the statistic.
Value Length	vInt	Contains the length of the value.
Value	string	Contains the statistic value.

The values *Name Length*, *Name*, *Value Length* and *Value* recur for each statistic requested.

[Report a bug](#)

## 13.7. HOT ROD OPERATION VALUES

The following is a list of valid *opcode* values for a request header and their corresponding response header values:

**Table 13.36. Opcode Request and Response Header Values**

Operation	Request Operation Code	Response Operation Code
put	0x01	0x02
get	0x03	0x04
putIfAbsent	0x05	0x06

Operation	Request Operation Code	Response Operation Code
replace	0x07	0x08
replaceIfUnmodified	0x09	0x0A
remove	0x0B	0x0C
removeIfUnmodified	0x0D	0x0E
containsKey	0x0F	0x10
clear	0x13	0x14
stats	0x15	0x16
ping	0x17	0x18
bulkGet	0x19	0x1A
getWithMetadata	0x1B	0x1C
bulkKeysGet	0x1D	0x1E
query	0x1F	0x20

Additionally, if the response header *opcode* value is **0x50**, it indicates an error response.

[Report a bug](#)

### 13.7.1. Magic Values

The following is a list of valid values for the *Magic* field in request and response headers:

**Table 13.37. Magic Field Values**

Value	Details
0xA0	Cache request marker.
0xA1	Cache response marker.

[Report a bug](#)

### 13.7.2. Status Values

The following is a table that contains all valid values for the *Status* field in a response header:

**Table 13.38. Status Values**

Value	Details
0x00	No error.
0x01	Not put/removed/replaced.
0x02	Key does not exist.
0x81	Invalid Magic value or Message ID.
0x82	Unknown command.
0x83	Unknown version.
0x84	Request parsing error.
0x85	Server error.
0x86	Command timed out.

[Report a bug](#)

### 13.7.3. Transaction Type Values

The following is a list of valid values for *Transaction Type* in a request header:

**Table 13.39. Transaction Type Field Values**

Value	Details
0	Indicates a non-transactional call or that the client does not support transactions. If used, the <b>TX_ID</b> field is omitted.
1	Indicates X/Open XA transaction ID (XID). This value is currently not supported.

[Report a bug](#)

### 13.7.4. Client Intelligence Values

The following is a list of valid values for *Client Intelligence* in a request header:

**Table 13.40. Client Intelligence Field Values**

Value	Details
0x01	Indicates a basic client that does not require any cluster or hash information.
0x02	Indicates a client that is aware of topology and requires cluster information.
0x03	Indicates a client that is aware of hash and distribution and requires both the cluster and hash information.

[Report a bug](#)

### 13.7.5. Flag Values

The following is a list of valid *flag* values in the request header:

**Table 13.41. Flag Field Values**

Value	Details
0x0001	ForceReturnPreviousValue

[Report a bug](#)

### 13.7.6. Hot Rod Error Handling

**Table 13.42. Hot Rod Error Handling using Response Header Fields**

Field	Data Type	Details
Error Opcode	-	Contains the error operation code.
Error Status Number	-	Contains a status number that corresponds to the <i><b>error opcode</b></i> .
Error Message Length	vInt	Contains the length of the error message.
Error Message	string	Contains the actual error message. If an <b>0x84</b> error code returns, which indicates that there was an error in parsing the request, this field contains the latest version supported by the <b>Hot Rod</b> server.



[Report a bug](#)

## 13.8. PUT REQUEST EXAMPLE

The following is the coded request from a sample `put` request using Hot Rod:

**Table 13.43. Put Request Example**

Byte	0	1	2	3	4	5	6	7
8	0xA0	0x09	0x41	0x01	0x07	0x4D ('M')	0x79 ('y')	0x43 ('C')
16	0x61 ('a')	0x63 ('c')	0x68 ('h')	0x65 ('e')	0x00	0x03	0x00	0x00
24	0x00	0x05	0x48 ('H')	0x65 ('e')	0x6C ('l')	0x6C ('l')	0x6F ('o')	0x00
32	0x00	0x05	0x57 ('W')	0x6F ('o')	0x72 ('r')	0x6C ('l')	0x64 ('d')	-

The following table contains all header fields and their values for the example request:

**Table 13.44. Example Request Field Names and Values**

Field Name	Byte	Value
Magic	0	0xA0
Version	2	0x41
Cache Name Length	4	0x07
Flag	12	0x00
Topology ID	14	0x00
Transaction ID	16	0x00
Key	18-22	'Hello'
Max Idle	24	0x00
Value	26-30	'World'
Message ID	1	0x09

Field Name	Byte	Value
Opcode	3	0x01
Cache Name	5-11	'MyCache'
Client Intelligence	13	0x03
Transaction Type	15	0x00
Key Field Length	17	0x05
Lifespan	23	0x00
Value Field Length	25	0x05

The following is a coded response for the sample **put** request:

**Table 13.45. Coded Response for the Sample Put Request**

Byte	0	1	2	3	4	5	6	7
8	0xA1	0x09	0x01	0x00	0x00	-	-	-

The following table contains all header fields and their values for the example response:

**Table 13.46. Example Response Field Names and Values**

Field Name	Byte	Value
Magic	0	0xA1
Opcode	2	0x01
Topology Change Marker	4	0x00
Message ID	1	0x09
Status	3	0x00

[Report a bug](#)

## 13.9. HOT ROD JAVA CLIENT

Hot Rod is a binary, language neutral protocol. A Java client is able to interact with a server via the Hot Rod protocol using the Hot Rod Java Client API.

[Report a bug](#)

### 13.9.1. Hot Rod Java Client Download

Use the following steps to download the JBoss Data Grid Hot Rod Java Client:

#### Procedure 13.2. Download Hot Rod Java Client

1. Log into the Customer Portal at <https://access.redhat.com>.
2. Click the **Downloads** button near the top of the page.
3. In the **Product Downloads** page, click **Red Hat JBoss Data Grid**.
4. Select the appropriate JBoss Data Grid version from the **Version:** drop down menu.
5. Locate the **Red Hat JBoss Data Grid \${VERSION} Hot Rod Java Client** entry and click the corresponding **Download** link.

[Report a bug](#)

### 13.9.2. Hot Rod Java Client Configuration

The Hot Rod Java client is configured both programmatically and externally using a configuration file or a properties file. The following example illustrate creation of a client instance using the available Java fluent API:

#### Example 13.1. Client Instance Creation

```
org.infinispan.client.hotrod.configuration.ConfigurationBuilder cb
= new org.infinispan.client.hotrod.configuration.ConfigurationBuilder();
cb.tcpNoDelay(true)
  .connectionPool()
    .numTestsPerEvictionRun(3)
    .testOnBorrow(false)
    .testOnReturn(false)
    .testWhileIdle(true)
  .addServer()
    .host("localhost")
    .port(11222);
RemoteCacheManager rmc = new RemoteCacheManager(cb.build());
```

#### Configuring the Hot Rod Java client using a properties file

To configure the Hot Rod Java client, edit the `hotrod-client.properties` file on the classpath.

The following example shows the possible content of the `hotrod-client.properties` file.

#### Example 13.2. Configuration

```
infinispan.client.hotrod.transport_factory =
org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory

infinispan.client.hotrod.server_list = 127.0.0.1:11222
```

```
infinispan.client.hotrod.marshaller =
org.infinispan.commons.marshall.jboss.GenericJBossMarshaller

infinispan.client.hotrod.async_executor_factory =
org.infinispan.client.hotrod.impl.async.DefaultAsyncExecutorFactory

infinispan.client.hotrod.default_executor_factory.pool_size = 1

infinispan.client.hotrod.default_executor_factory.queue_size = 10000

infinispan.client.hotrod.hash_function_impl.1 =
org.infinispan.client.hotrod.impl.consistenthash.ConsistentHashV1

infinispan.client.hotrod.tcp_no_delay = true

infinispan.client.hotrod.ping_on_startup = true

infinispan.client.hotrod.request_balancing_strategy =
org.infinispan.client.hotrod.impl.transport.tcp.RoundRobinBalancingStrategy

infinispan.client.hotrod.key_size_estimate = 64

infinispan.client.hotrod.value_size_estimate = 512

infinispan.client.hotrod.force_return_values = false

infinispan.client.hotrod.tcp_keep_alive = true

## below is connection pooling config

maxActive=-1

maxTotal = -1

maxIdle = -1


whenExhaustedAction = 1

timeBetweenEvictionRunsMillis=120000

minEvictableIdleTimeMillis=300000

testWhileIdle = true

minIdle = 1
```



## NOTE

The **TCP KEEPALIVE** configuration is enabled/disabled on the Hot Rod Java client either through a config property as seen in the example (`infinispan.client.hotrod.tcp_keep_alive = true/false` or programmatically through the `org.infinispan.client.hotrod.ConfigurationBuilder.tcpKeepAlive()` method.

Either of the following two constructors must be used in order for the properties file to be consumed by Red Hat JBoss Data Grid:

1. `new RemoteCacheManager(boolean start)`
2. `new RemoteCacheManager()`

[Report a bug](#)

### 13.9.3. Hot Rod Java Client Basic API

The following code shows how the client API can be used to store or retrieve information from a Hot Rod server using the Hot Rod Java client. This example assumes that a Hot Rod server has been started bound to the default location, `localhost:11222`.

#### Example 13.3. Basic API

```
//API entry point, by default it connects to localhost:11222
    BasicCacheContainer cacheContainer = new RemoteCacheManager();
//obtain a handle to the remote default cache
    BasicCache<String, String> cache = cacheContainer.getCache();
//now add something to the cache and ensure it is there
    cache.put("car", "ferrari");
    assert cache.get("car").equals("ferrari");
//remove the data
    cache.remove("car");
    assert !cache.containsKey("car") : "Value must have been
removed!";
```

The **RemoteCacheManager** corresponds to **DefaultCacheManager**, and both implement **BasicCacheContainer**.

This API facilitates migration from local calls to remote calls via Hot Rod. This can be done by switching between **DefaultCacheManager** and **RemoteCacheManager**, which is simplified by the common **BasicCacheContainer** interface.

All keys can be retrieved from the remote cache using the **keySet()** method. If the remote cache is a distributed cache, the server will start a Map/Reduce job to retrieve all keys from clustered nodes and return all keys to the client.

Use this method with caution if there are a large number of keys.

```
Set keys = remoteCache.keySet();
```

[Report a bug](#)

### 13.9.4. Hot Rod Java Client Versioned API

To ensure data consistency, Hot Rod stores a version number that uniquely identifies each modification. Using *getVersioned*, clients can retrieve the value associated with the key as well as the current version.

When using the Hot Rod Java client, a **RemoteCacheManager** provides instances of the **RemoteCache** interface that accesses the named or default cache on the remote cluster. This extends the **Cache** interface to which it adds new methods, including the versioned API.

#### Example 13.4. Using Versioned Methods

```
// To use the versioned API, remote classes are specifically needed
RemoteCacheManager remoteCacheManager = new RemoteCacheManager();
RemoteCache<String, String> remoteCache = remoteCacheManager.getCache();
remoteCache.put("car", "ferrari");
VersionedValue valueBinary = remoteCache.getVersioned("car");
// removal only takes place only if the version has not been changed
// in between. (a new version is associated with 'car' key on each
change)
assert remoteCache.removeWithVersion("car", valueBinary.getVersion());
assert !remoteCache.containsKey("car");
```

#### Example 13.5. Using Replace

```
remoteCache.put("car", "ferrari");
VersionedValue valueBinary = remoteCache.getVersioned("car");
assert remoteCache.replaceWithVersion("car", "lamborghini",
valueBinary.getVersion());
```

[Report a bug](#)

## 13.10. HOT ROD C++ CLIENT

The Hot Rod C++ client is a new addition to the Hot Rod client family which includes the Hot Rod Java client. It enables C++ runtime applications to connect and interact with Red Hat JBoss Data Grid remote servers.

The Hot Rod C++ client allows applications developed in C++ to read or write data to remote caches. The Hot Rod C++ client supports all three levels of client intelligence and is supported on the following platforms:

- Red Hat Enterprise Linux 5, 64-bit
- Red Hat Enterprise Linux 6, 64-bit
- Red Hat Enterprise Linux 7, 64-bit

The Hot Rod C++ client is available as a Technology Preview on 64-bit Windows with Visual Studio 2010.

[Report a bug](#)

### 13.10.1. Hot Rod C++ Client Formats

The Hot Rod C++ client is available in the following two library formats:

- Static library
- Shared/Dynamic library

#### Static Library

The static library is statically linked to an application. This increases the size of the final executable. The application is self-contained and it does not need to ship a separate library.

#### Shared/Dynamic Library

Shared/Dynamic libraries are dynamically linked to an application at runtime. The library is stored in a separate file and can be upgraded separately from the application, without recompiling the application.



#### NOTE

This can only happen if the library's major version is equal to the one against which the application was linked at compile time, indicating that it is binary compatible.

[Report a bug](#)

### 13.10.2. Hot Rod C++ Client Prerequisites

In order to use the Hot Rod C++ Client, the following are needed:

- C++ 03 compiler with support for shared\_ptr TR1 (GCC 4.0+, Visual Studio C++ 2010).
- Red Hat JBoss Data Grid Server 6.1.0 or higher version.

[Report a bug](#)

### 13.10.3. Hot Rod C++ Client Download

The Hot Rod C++ client is included in a separate zip file **jboss-datagrid-<version>-hotrod-cpp-client-<platform>.zip** under Red Hat JBoss Data Grid binaries on the Red Hat Customer Portal at <https://access.redhat.com>. Download the appropriate Hot Rod C++ client which applies to your operating system.

[Report a bug](#)

### 13.10.4. Hot Rod C++ Client Configuration

The Hot Rod C++ client interacts with a remote Hot Rod server using the RemoteCache API. To initiate communication with a particular Hot Rod server, configure RemoteCache and choose the specific cache on the Hot Rod server.

Use the ConfigurationBuilder API to configure:

- The initial set of servers to connect to.

- Connection pooling attributes.
- Connection/Socket timeouts and TCP nodelay.
- Hot Rod protocol version.

### Sample C++ main executable file configuration

The following example shows how to use the ConfigurationBuilder to configure a RemoteCacheManager and how to obtain the default remote cache:

#### Example 13.6. SimpleMain.cpp

```
#include "infinispan/hotrod/ConfigurationBuilder.h"
#include "infinispan/hotrod/RemoteCacheManager.h"
#include "infinispan/hotrod/RemoteCache.h"
#include <stdlib.h>
using namespace infinispan::hotrod;
int main(int argc, char** argv) {
    ConfigurationBuilder b;
    b.addServer().host("127.0.0.1").port(11222);
    RemoteCacheManager cm(builder.build());
    RemoteCache<std::string, std::string> cache =
cm.getCache<std::string, std::string>();
    return 0;
}
```

[Report a bug](#)

### 13.10.5. Hot Rod C++ Client API

The RemoteCacheManager is a starting point to obtain a reference to a RemoteCache. The RemoteCache API can interact with a remote Hot Rod server and the specific cache on that server.

Using the RemoteCache reference obtained in the previous example, it is possible to put, get, replace and remove values in a remote cache. It is also possible to perform bulk operations, such as retrieving all of the keys, and clearing the cache.

When a RemoteCacheManager is stopped, all resources in use are released.

#### Example 13.7. SimpleMain.cpp

```
RemoteCache<std::string, std::string> rc = cm.getCache<std::string,
std::string>();
    std::string k1("key13");
    std::string v1("boron");
    // put
    rc.put(k1, v1);
    std::auto_ptr<std::string> rv(rc.get(k1));
    rc.putIfAbsent(k1, v1);
    std::auto_ptr<std::string> rv2(rc.get(k1));
    std::map<HR_SHARED_PTR<std::string>, HR_SHARED_PTR<std::string> > map
= rc.getBulk(0);
    std::cout << "getBulk size" << map.size() << std::endl;
```



```
..  
.cm.stop();
```

[Report a bug](#)

## 13.11. HOT ROD C# CLIENT

The Hot Rod C# client is a new addition to the list of Hot Rod clients that includes Hot Rod Java and Hot Rod C++ clients. Hot Rod C# client allows .NET runtime applications to connect and interact with Red Hat JBoss Data Grid servers.

The Hot Rod C# client is aware of the cluster topology and hashing scheme, and can access an entry on the server in a single hop similar to the Hot Rod Java and Hot Rod C++ clients.

The Hot Rod C# client is compatible with 32-bit and 64-bit operating systems on which the .NET Framework is supported by Microsoft. The .NET Framework 4.0 is a prerequisite along with the supported operating systems to use the Hot Rod C# client.

[Report a bug](#)

### 13.11.1. Hot Rod C# Client Download and Installation

The Hot Rod C# client is included in a .msi file **jboss-datagrid-<version>-hotrod-dotnet-client.msi** packed for download with Red Hat JBoss Data Grid . To install the Hot Rod C# client, execute the following instructions.

#### Procedure 13.3. Installing the Hot Rod C# Client

1. As an administrator, navigate to the location where the Hot Rod C# .msi file is downloaded. Run the .msi file to launch the windows installer and then click **Next**.




Figure 13.1. Hot Rod C# Client Setup Welcome

2. Review the end-user license agreement. Select the **I accept the terms in the License Agreement** check box and then click **Next**.


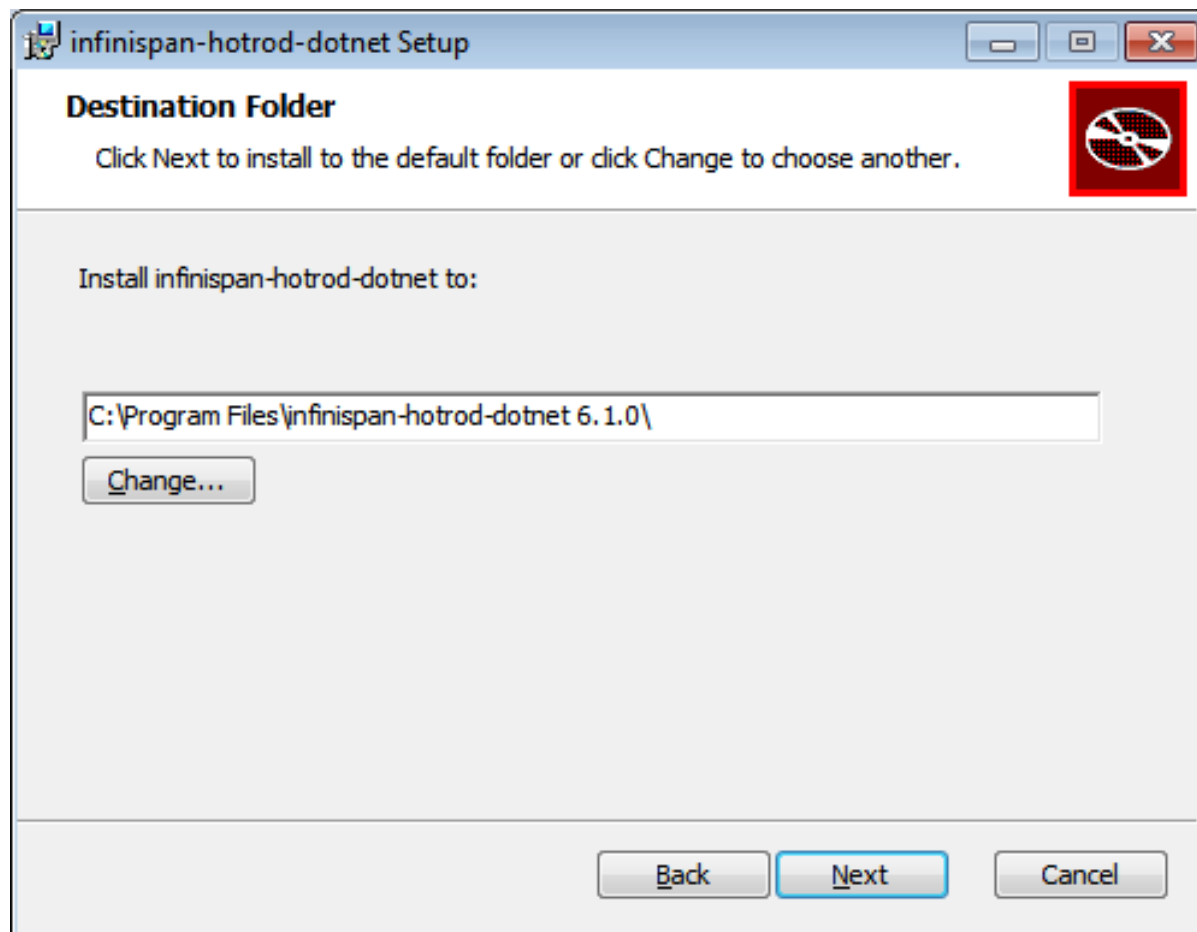


Figure 13.2. Hot Rod C# Client End-User License Agreement

3. To change the default directory, click **Change...** or click **Next** to install in the default directory.



**Figure 13.3. Hot Rod C# Client Destination Folder**

4. Click **Finish** to complete the Hot Rod C# client installation.




Figure 13.4. Hot Rod C# Client Setup Completion

[Report a bug](#)

### 13.11.2. Hot Rod C# Client Configuration

The Hot Rod C# client is configured programmatically using the `ConfigurationBuilder`. Configure the host and the port to which the client should connect.

#### Sample C# file configuration

The following example shows how to use the `ConfigurationBuilder` to configure a `RemoteCacheManager`.

#### Example 13.8. C# configuration

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Threading.Tasks;
using Infinispan.HotRod;
using Infinispan.HotRod.Config;
namespace simpleapp
{
    class Program
    {
        static void Main(string[] args)
        {
```

```

        ConfigurationBuilder builder = new ConfigurationBuilder();
        builder.AddServer()
            .Host(args.Length > 1 ? args[0] : "127.0.0.1")
            .Port(args.Length > 2 ? int.Parse(args[1]) : 11222);
        Configuration config = builder.Build();
        RemoteCacheManager cacheManager = new
RemoteCacheManager(config);
        [...]
    }
}
}

```

[Report a bug](#)

### 13.11.3. Hot Rod C# Client API

The **RemoteCacheManager** is a starting point to obtain a reference to a **RemoteCache**.

The following example shows retrieval of a default cache from the server and a few basic operations.

#### Example 13.9.

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Threading.Tasks;
using Infinispan.HotRod;
using Infinispan.HotRod.Config;
namespace simpleapp
{
    class Program
    {
        static void Main(string[] args)
        {
            ConfigurationBuilder builder = new ConfigurationBuilder();
            builder.AddServer()
                .Host(args.Length > 1 ? args[0] : "127.0.0.1")
                .Port(args.Length > 2 ? int.Parse(args[1]) : 11222);
            Configuration config = builder.Build();
            RemoteCacheManager cacheManager = new
RemoteCacheManager(config);
            cacheManager.Start();
            // Retrieve a reference to the default cache.
            IRemoteCache<String, String> cache =
cacheManager.GetCache<String, String>();
            // Add entries.
            cache.Put("key1", "value1");
            cache.PutIfAbsent("key1", "anotherValue1");
            cache.PutIfAbsent("key2", "value2");
            cache.PutIfAbsent("key3", "value3");
            // Retrieve entries.
            Console.WriteLine("key1 -> " + cache.Get("key1"));
            // Bulk retrieve key/value pairs.

```

```

        int limit = 10;
        IDictionary<String, String> result = cache.GetBulk(limit);
        foreach (KeyValuePair<String, String> kv in result)
        {
            Console.WriteLine(kv.Key + " -> " + kv.Value);
        }
        // Remove entries.
        cache.Remove("key2");
        Console.WriteLine("key2 -> " + cache.Get("key2"));
        cacheManager.Stop();
    }
}

```

[Report a bug](#)

### 13.11.4. String Marshaller for Interoperability

To use the string compatibility marshaller, enable the compatibility mode on the server-side. On the C# client-side, pass an instance of `CompatibilitySerializer` to the `RemoteCacheManager` constructor similar to this:

```

[... ]
RemoteCacheManager cacheManager = new RemoteCacheManager(new
CompatibilitySerializer());
[... ]
cache.Put("key", "value");
String value = cache.Get("key");
[... ]

```



#### NOTE

Attempts to store or retrieve non-string key/values will result in a `HotRodClientException` being thrown.

[Report a bug](#)

## 13.12. INTEROPERABILITY BETWEEN HOT ROD C++ AND HOT ROD JAVA CLIENT

Red Hat JBoss Data Grid provides interoperability between Hot Rod Java and Hot Rod C++ clients to access structured data. This is made possible by structuring and serializing data using Google's Protobuf format.

For example, using interoperability between languages would allow a Hot Rod C++ client to write the following `Person` object structured and serialized using Protobuf, and the Java C++ client can read the same `Person` object structured as Protobuf.

### Example 13.10. Using Interoperability Between Languages

```

package sample;

```

```
message Person {  
    required int32 age = 1;  
    required string name = 2;  
}
```

Although it is neither enforced or supported, it is recommended that the Hot Rod Java client use the shipped Protostream library for serializing Protobuf objects. Protobuf library is a recommended serialization solution for C++ (<https://code.google.com/p/protobuf/>).


Interoperability between C++ and Hot Rod Java Client is fully supported for primitive data types, strings, and byte arrays, as Protobuf and Protostream are not required for these types of interoperability.

For information about how the C++ client can read Protobuf encoded data, see <https://github.com/jboss-developer/jboss-jdg-quickstarts/tree/master/remote-query/src/main/cpp>.

[Report a bug](#)

## 13.13. COMPATIBILITY BETWEEN SERVER AND HOT ROD CLIENT VERSIONS

Hot Rod clients, such as the Hot Rod Java, Hot Rod C++, and Hot Rod C#, are compatible with different versions of Red Hat JBoss Data Grid server. The server should be of the latest version in order to run with different Hot Rod clients.



### NOTE

It is recommended to use the same version of the Hot Rod client and the JBoss Data Grid server, except in a case of migration or upgrade, to prevent any known problems.

Consider the following scenarios.

#### Scenario 1: Server running on a newer version than the Hot Rod client.

The following will be the impact on the client side:

- client will not have advantage of the latest protocol improvements.
- client might run into known issues which are fixed for the server-side version.
- client can only use the functionalities available in its current version and the previous versions.

#### Scenario 2: Hot Rod client running on a newer version than the server.

In this case, when a Hot Rod client connects to a JBoss Data Grid server, the connection will be rejected with an exception error. The client can be downgraded to a known protocol version by setting the client side property -

***org.infinispan.client.hotrod.configuration.ConfigurationBuilder.protocolVersion*** to a specific version, such as Hot Rod 1.3, as a String. In this case the client is able to connect to the server, but will be restricted to the functionality of that version. Any command which is not supported by this protocol version will not work and throw an exception. Also, the topology information might be inefficient in this case.



**NOTE**

It is not recommended to use this approach without guidance from Red Hat support.

The following table details the compatibility between different Hot Rod client and server versions.

**Table 13.47. Hot Rod client and server compatibility**

JBoss Data Grid Server Version	Hot Rod Protocol Version
JBoss Data Grid 6.5.1	Hot Rod 2.2
JBoss Data Grid 6.5.0	Hot Rod 2.2
JBoss Data Grid 6.4.1	Hot Rod 2.0
JBoss Data Grid 6.4.0	Hot Rod 2.0
JBoss Data Grid 6.3.2	Hot Rod 2.0
JBoss Data Grid 6.3.1	Hot Rod 2.0
JBoss Data Grid 6.3.0	Hot Rod 2.0
JBoss Data Grid 6.2.1	Hot Rod 1.3
JBoss Data Grid 6.2.0	Hot Rod 1.3
JBoss Data Grid 6.1.0	Hot Rod 1.2
JBoss Data Grid 6.0.1	Hot Rod 1.1
JBoss Data Grid 6.0.0	Hot Rod 1.1

[Report a bug](#)

## **PART VI. SET UP LOCKING FOR THE CACHE**

## CHAPTER 14. LOCKING

Red Hat JBoss Data Grid provides locking mechanisms to prevent dirty reads (where a transaction reads an outdated value before another transaction has applied changes to it) and non-repeatable reads.

[Report a bug](#)

### 14.1. CONFIGURE LOCKING (REMOTE CLIENT-SERVER MODE)

In Remote Client-Server mode, locking is configured using the **locking** element within the cache tags (for example, **invalidation-cache**, **distributed-cache**, **replicated-cache** or **local-cache**).



#### NOTE

The default isolation mode for the Remote Client-Server mode configuration is **READ\_COMMITTED**. If the **isolation** attribute is included to explicitly specify an isolation mode, it is ignored, a warning is thrown, and the default value is used instead.

The following is a sample procedure of a basic locking configuration for a default cache in Red Hat JBoss Data Grid's Remote Client-Server mode.

#### Procedure 14.1. Configure Locking (Remote Client-Server Mode)

```
<distributed-cache>
  <locking acquire-timeout="30000"
    concurrency-level="1000"
    striping="false" />
  <!-- Additional configuration here -->
</distributed-cache>
```

1. The **acquire-timeout** parameter specifies the number of milliseconds after which lock acquisition will time out.
2. The **concurrency-level** parameter defines the number of lock stripes used by the LockManager.
3. The **striping** parameter specifies whether lock striping will be used for the local cache.

[Report a bug](#)

### 14.2. CONFIGURE LOCKING (LIBRARY MODE)

For Library mode, the **locking** element and its parameters are set within the **default** element and for each named cache, it occurs within the **namedCache** element. The following is an example of this configuration:

#### Procedure 14.2. Configure Locking (Library Mode)

```
<infinispan>
  <!-- Other configuration elements here -->
  <default>
```

```
<locking concurrencyLevel="${VALUE}"
  isolationLevel="${LEVEL}"
  lockAcquisitionTimeout="${TIME}"
  useLockStriping="${TRUE/FALSE}"
  writeSkewCheck="${TRUE/FALSE}" />
```

1. The **concurrencyLevel** parameter specifies the concurrency level for the lock container. Set this value according to the number of concurrent threads interacting with the data grid.
2. The **isolationLevel** parameter specifies the cache's isolation level. Valid isolation levels are **READ\_COMMITTED** and **REPEATABLE\_READ**. For details about isolation levels, see [Section 16.1, "About Isolation Levels"](#)
3. The **lockAcquisitionTimeout** parameter specifies time (in milliseconds) after which a lock acquisition attempt times out.
4. The **useLockStriping** parameter specifies whether a pool of shared locks are maintained for all entries that require locks. If set to **FALSE**, locks are created for each entry in the cache. For details, see [Section 15.1, "About Lock Striping"](#)
5. The **writeSkewCheck** parameter is only valid if the **isolationLevel** is set to **REPEATABLE\_READ**. If this parameter is set to **FALSE**, a disparity between a working entry and the underlying entry at write time results in the working entry overwriting the underlying entry. If the parameter is set to **TRUE**, such conflicts (namely write skews) throw an exception. The **writeSkewCheck** parameter can be only used with **OPTIMISTIC** transactions and it requires entry versioning to be enabled, with **SIMPLE** versioning scheme.

[Report a bug](#)

## 14.3. LOCKING TYPES

### 14.3.1. About Optimistic Locking

Optimistic locking allows multiple transactions to complete simultaneously by deferring lock acquisition to the transaction prepare time.

Optimistic mode assumes that multiple transactions can complete without conflict. It is ideal where there is little contention between multiple transactions running concurrently, as transactions can commit without waiting for other transaction locks to clear. With **writeSkewCheck** enabled, transactions in optimistic locking mode roll back if one or more conflicting modifications are made to the data before the transaction completes.

[Report a bug](#)

### 14.3.2. About Pessimistic Locking

Pessimistic locking is also known as eager locking.

Pessimistic locking prevents more than one transaction to modify a value of a key by enforcing cluster-wide locks on each write operation. Locks are only released once the transaction is completed either through committing or being rolled back.

Pessimistic mode is used where a high contention on keys is occurring, resulting in inefficiencies and unexpected roll back operations.

[Report a bug](#)

### 14.3.3. Pessimistic Locking Types

Red Hat JBoss Data Grid includes explicit pessimistic locking and implicit pessimistic locking:

- Explicit Pessimistic Locking, which uses the JBoss Data Grid Lock API to allow cache users to explicitly lock cache keys for the duration of a transaction. The Lock call attempts to obtain locks on specified cache keys across all nodes in a cluster. This attempt either fails or succeeds for all specified cache keys. All locks are released during the commit or rollback phase.
- Implicit Pessimistic Locking ensures that cache keys are locked in the background as they are accessed for modification operations. Using Implicit Pessimistic Locking causes JBoss Data Grid to check and ensure that cache keys are locked locally for each modification operation. Discovering unlocked cache keys causes JBoss Data Grid to request a cluster-wide lock to acquire a lock on the unlocked cache key.

[Report a bug](#)

### 14.3.4. Explicit Pessimistic Locking Example

The following is an example of explicit pessimistic locking that depicts a transaction that runs on one of the cache nodes:

#### Procedure 14.3. Transaction with Explicit Pessimistic Locking

```
tx.begin()
cache.lock(K)
cache.put(K, V5)
tx.commit()
```

1. When the line `cache.lock(K)` executes, a cluster-wide lock is acquired on **K**.
2. When the line `cache.put(K, V5)` executes, it guarantees success.
3. When the line `tx.commit()` executes, the locks held for this process are released.

[Report a bug](#)

### 14.3.5. Implicit Pessimistic Locking Example

An example of implicit pessimistic locking using a transaction that runs on one of the cache nodes is as follows:

#### Procedure 14.4. Transaction with Implicit Pessimistic locking

```
tx.begin()
cache.put(K, V)
cache.put(K2, V2)
cache.put(K, V5)
tx.commit()
```

1. When the line `cache.put(K, V)` executes, a cluster-wide lock is acquired on **K**.

2. When the line `cache.put(K2, V2)` executes, a cluster-wide lock is acquired on **K2**.
3. When the line `cache.put(K, V5)` executes, the lock acquisition is non operational because a cluster-wide lock for **K** has been previously acquired. The `put` operation will still occur.
4. When the line `tx.commit()` executes, all locks held for this transaction are released.

[Report a bug](#)

### 14.3.6. Configure Locking Mode (Remote Client-Server Mode)

To configure a locking mode in Red Hat JBoss Data Grid's Remote Client-Server mode, use the `transaction` element as follows:

```
<transaction locking="{OPTIMISTIC/PESSIMISTIC}" />
```

[Report a bug](#)

### 14.3.7. Configure Locking Mode (Library Mode)

In Red Hat JBoss Data Grid's Library mode, the locking mode is set within the `transaction` element as follows:

```
<transaction transactionManagerLookupClass="
{TransactionManagerLookupClass}"
    transactionMode="{TRANSACTIONAL, NON_TRANSACTIONAL}"
    lockingMode="{OPTIMISTIC, PESSIMISTIC}"
    useSynchronization="true">
</transaction>
```

Set the `lockingMode` value to **OPTIMISTIC** or **PESSIMISTIC** to configure the locking mode used for the transactional cache.

[Report a bug](#)

## 14.4. LOCKING OPERATIONS

### 14.4.1. About the LockManager

The **LockManager** component is responsible for locking an entry before a write process initiates. The **LockManager** uses a **LockContainer** to locate, hold and create locks. There are two types of **LockContainers** JBoss Data Grid uses internally and their choice is dependent on the `useLockStriping` setting. The first type offers support for lock striping while the second type supports one lock per entry.

**See Also:**

- [Chapter 15, Set Up Lock Striping](#)

[Report a bug](#)

### 14.4.2. About Lock Acquisition

Red Hat JBoss Data Grid acquires remote locks lazily by default. The node running a transaction locally acquires the lock while other cluster nodes attempt to lock cache keys that are involved in a two phase prepare/commit phase. JBoss Data Grid can lock cache keys in a pessimistic manner either explicitly or implicitly.

[Report a bug](#)

### 14.4.3. About Concurrency Levels

Concurrency refers to the number of threads simultaneously interacting with the data grid. In Red Hat JBoss Data Grid, concurrency levels refer to the number of concurrent threads used within a lock container.

In JBoss Data Grid, concurrency levels determine the size of each striped lock container. Additionally, concurrency levels tune all related JDK **ConcurrentHashMap** based collections, such as those internal to **DataContainers**.

[Report a bug](#)

## CHAPTER 15. SET UP LOCK STRIPING

### 15.1. ABOUT LOCK STRIPING

Lock Striping allocates locks from a shared collection of (fixed size) locks in the cache. Lock allocation is based on the hash code for each entry's key. Lock Striping provides a highly scalable locking mechanism with fixed overhead. However, this comes at a cost of potentially unrelated entries being blocked by the same lock.

Lock Striping is disabled by default in Red Hat JBoss Data Grid. If lock striping remains disabled, a new lock is created for each entry. This alternative approach can provide greater concurrent throughput, but also results in additional memory usage, garbage collection churn, and other disadvantages.

[Report a bug](#)

### 15.2. CONFIGURE LOCK STRIPING (REMOTE CLIENT-SERVER MODE)

Lock striping in Red Hat JBoss Data Grid's Remote Client-Server mode is enabled by setting the *striping* element to **true**.

#### Example 15.1. Lock Striping (Remote Client-Server Mode)

```
<locking acquire-timeout="20000"  
  concurrency-level="500"  
  striping="true" />
```



#### NOTE

The default isolation mode for the Remote Client-Server mode configuration is **READ\_COMMITTED**. If the *isolation* attribute is included to explicitly specify an isolation mode, it is ignored, a warning is thrown, and the default value is used instead.

The *locking* element uses the following attributes:

- The *acquire-timeout* attribute specifies the maximum time to attempt a lock acquisition. The default value for this attribute is **10000** milliseconds.
- The *concurrency-level* attribute specifies the concurrency level for lock containers. Adjust this value according to the number of concurrent threads interacting with JBoss Data Grid. The default value for this attribute is **32**.
- The *striping* attribute specifies whether a shared pool of locks is maintained for all entries that require locking (**true**). If set to **false**, a lock is created for each entry. Lock striping controls the memory footprint but can reduce concurrency in the system. The default value for this attribute is **false**.

[Report a bug](#)

### 15.3. CONFIGURE LOCK STRIPING (LIBRARY MODE)



Lock striping is disabled by default in Red Hat JBoss Data Grid. Configure lock striping in JBoss Data Grid's Library mode using the *useLockStriping* parameter as demonstrated in the following procedure.

### Procedure 15.1. Configure Lock Striping (Library Mode)

```
<infinispan>
  <!-- Additional configuration elements here -->
  <default>

    <locking concurrencyLevel="${VALUE}"
      isolationLevel="${LEVEL}"
      lockAcquisitionTimeout="${TIME}"
      useLockStriping="${TRUE/FALSE}"
      writeSkewCheck="${TRUE/FALSE}" />
    <!-- Additional configuration elements here -->
  </default>
</infinispan>
```

1. The *concurrencyLevel* is used to specify the size of the shared lock collection use when lock striping is enabled.
2. The *isolationLevel* parameter specifies the cache's isolation level. Valid isolation levels are **READ\_COMMITTED** and **REPEATABLE\_READ**.
3. The *lockAcquisitionTimeout* parameter specifies time (in milliseconds) after which a lock acquisition attempt times out.
4. The *useLockStriping* parameter specifies whether a pool of shared locks are maintained for all entries that require locks. If set to **FALSE**, locks are created for each entry in the cache. If set to **TRUE**, lock striping is enabled and shared locks are used as required from the pool.
5. The *writeSkewCheck* check determines if a modification to the entry from a different transaction should roll back the transaction. Write skew set to true requires *isolation\_level* set to **REPEATABLE\_READ**. The default value for *writeSkewCheck* and *isolation\_level* are **FALSE** and **READ\_COMMITTED** respectively. The *writeSkewCheck* parameter can be only used with **OPTIMISTIC** transactions and it requires entry versioning to be enabled, with **SIMPLE** versioning scheme.

[Report a bug](#)

## CHAPTER 16. SET UP ISOLATION LEVELS

### 16.1. ABOUT ISOLATION LEVELS

Isolation levels determine when readers can view a concurrent write. **READ\_COMMITTED** and **REPEATABLE\_READ** are the two isolation modes offered in Red Hat JBoss Data Grid.

- **READ\_COMMITTED**. This isolation level is applicable to a wide variety of requirements. This is the default value in Remote Client-Server and Library modes.
- **REPEATABLE\_READ**.



#### IMPORTANT

The only valid value for locks in Remote Client-Server mode is the default **READ\_COMMITTED** value. The value explicitly specified with the *isolation* value is ignored.

If the **locking** element is not present in the configuration, the default isolation value is **READ\_COMMITTED**.

For isolation mode configuration examples in JBoss Data Grid, see the lock striping configuration samples:

- See [Section 15.2, “Configure Lock Striping \(Remote Client-Server Mode\)”](#) for a Remote Client-Server mode configuration sample.
- See [Section 15.3, “Configure Lock Striping \(Library Mode\)”](#) for a Library mode configuration sample.

[Report a bug](#)

### 16.2. ABOUT READ\_COMMITTED

**READ\_COMMITTED** is one of two isolation modes available in Red Hat JBoss Data Grid.

In JBoss Data Grid's **READ\_COMMITTED** mode, write operations are made to copies of data rather than the data itself. A write operation blocks other data from being written, however writes do not block read operations. As a result, both **READ\_COMMITTED** and **REPEATABLE\_READ** modes permit read operations at any time, regardless of when write operations occur.

In **READ\_COMMITTED** mode multiple reads of the same key within a transaction can return different results due to write operations in different transactions modifying data between reads. This phenomenon is known as non-repeatable reads and is avoided in **REPEATABLE\_READ** mode.

[Report a bug](#)

### 16.3. ABOUT REPEATABLE\_READ

**REPEATABLE\_READ** is one of two isolation modes available in Red Hat JBoss Data Grid.

Traditionally, **REPEATABLE\_READ** does not allow write operations while read operations are in progress, nor does it allow read operations when write operations occur. This prevents the "non-repeatable read"

phenomenon, which occurs when a single transaction has two read operations on the same row but the retrieved values differ (possibly due to a write operation modifying the value between the two read operations).

JBoss Data Grid's **REPEATABLE\_READ** isolation mode preserves the value of an entry before a modification occurs. As a result, the "non-repeatable read" phenomenon is avoided because a second read operation on the same entry retrieves the preserved value rather than the new modified value. As a result, the two values retrieved by the two read operations in a single transaction will always match, even if a write operation occurs in a different transaction between the two reads.

[Report a bug](#)

## **PART VII. SET UP AND CONFIGURE A CACHE STORE**

## CHAPTER 17. CACHE STORES

The cache store connects Red Hat JBoss Data Grid to the persistent data store. Cache stores are associated with individual caches. Different caches attached to the same cache manager can have different cache store configurations.



### NOTE

If a clustered cache is configured with an unshared cache store (where *shared* is set to **false**), on node join, stale entries which might have been removed from the cluster might still be present in the stores and can reappear.

[Report a bug](#)

### 17.1. CACHE LOADERS AND CACHE WRITERS

Integration with the persistent store is done through the following SPIs located in `org.infinispan.persistence.spi`:

- `CacheLoader`
- `CacheWriter`
- `AdvancedCacheLoader`
- `AdvancedCacheWriter`

`CacheLoader` and `CacheWriter` provide basic methods for reading and writing to a store. `CacheLoader` retrieves data from a data store when the required data is not present in the cache.

`AdvancedCacheLoader` and `AdvancedCacheWriter` provide operations to manipulate the underlying storage in bulk: parallel iteration and purging of expired entries, clear and size.

The `org.infinispan.persistence.file.SingleFileStore` is a good starting point to write your own store implementation.



### NOTE

Previously, JBoss Data Grid used the old API (`CacheLoader`, extended by `CacheStore`), which is also still available.

[Report a bug](#)

### 17.2. CACHE STORE CONFIGURATION

#### 17.2.1. Configuring the Cache Store

Cache stores can be configured in a chain. Cache read operations checks each cache store in the order configured until a valid non-null element of data has been located. Write operations affect all cache stores unless the *ignoreModifications* element has been set to **"true"** for a specific cache store.

[Report a bug](#)

## 17.2.2. Configure the Cache Store using XML (Library Mode)

The following example demonstrates cache store configuration using XML in JBoss Data Grid's Library mode:

```
<persistence passivation="false">
  <singleFile shared="false"
    preload="true"
    fetchPersistentState="true"
    ignoreModifications="false"
    purgeOnStartup="false"
    location="{java.io.tmpdir}" >
    <async enabled="true"
      flushLockTimeout="15000"
      threadPoolSize="5" />
    <singleton enabled="true"
      pushStateWhenCoordinator="true"
      pushStateTimeout="20000" />
  </singleFile>
</persistence>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, "Cache Store Configuration Details \(Library Mode\)"](#).

[Report a bug](#)

## 17.2.3. Configure the Cache Store Programmatically

The following example demonstrates how to configure the cache store programmatically:

```
ConfigurationBuilder builder = new ConfigurationBuilder();
builder.persistence()
  .passivation(false)
  .addSingleFileStore()
    .shared(false)
    .preload(true)
    .fetchPersistentState(true)
    .ignoreModifications(false)
    .purgeOnStartup(false)
    .location(System.getProperty("java.io.tmpdir"))
  .async()
    .enabled(true)
    .flushLockTimeout(15000)
    .threadPoolSize(5)
  .singleton()
    .enabled(true)
    .pushStateWhenCoordinator(true)
    .pushStateTimeout(20000);
```



### NOTE

This configuration is for a single-file cache store. Some attributes, such as *location* are specific to the single-file cache store and are not used for other types of cache stores.

### Procedure 17.1. Configure the Cache store Programmatically

1. Use the *ConfigurationBuilder* to create a new configuration object.
2. The *passivation* elements affects the way Red Hat JBoss Data Grid interacts with stores. Passivation removes an object from in-memory cache and writes it to a secondary data store, such as a system or database. Passivation is **false** by default.
3. The *addSingleFileStore()* elements adds the SingleFileStore as the cache store for this configuration. It is possible to create other stores, such as a JDBC Cache Store, which can be added using the *addStore* method.
4. The *shared* parameter indicates that the cache store is shared by different cache instances. For example, where all instances in a cluster use the same JDBC settings to talk to the same remote, shared database. *shared* is **false** by default. When set to **true**, it prevents duplicate data being written to the cache store by different cache instances.
5. The *preload* element is set to **false** by default. When set to **true** the data stored in the cache store is preloaded into the memory when the cache starts. This allows data in the cache store to be available immediately after startup and avoids cache operations delays as a result of loading data lazily. Preloaded data is only stored locally on the node, and there is no replication or distribution of the preloaded data. JBoss Data Grid will only preload up to the maximum configured number of entries in eviction.
6. The *fetchPersistentState* element determines whether or not to fetch the persistent state of a cache and apply it to the local cache store when joining the cluster. If the cache store is shared the fetch persistent state is ignored, as caches access the same cache store. A configuration exception will be thrown when starting the cache service if more than one cache store has this property set to **true**. The *fetchPersistentState* property is **false** by default.
7. The *ignoreModifications* element determines whether write methods are pushed to the specific cache store by allowing write operations to the local file cache store, but not the shared cache store. In some cases, transient application data should only reside in a file-based cache store on the same server as the in-memory cache. For example, this would apply with a further JDBC based cache store used by all servers in the network. *ignoreModifications* is **false** by default.
8. The *purgeOnStartup* element controls whether cache store is purged when it starts up and is **false** by default.
9. The *location* element configuration element sets a location on disk where the store can write.
10. These attributes configure aspects specific to each cache store. For example, the *location* attribute points to where the SingleFileStore will keep files containing data. Other stores may require more complex configuration.
11. The *singleton* element enables modifications to be stored by only one node in the cluster. This node is called the coordinator. The coordinator pushes the caches in-memory states to disk. This function is activated by setting the *enabled* attribute to **true** in all nodes. The *shared* parameter cannot be defined with *singleton* enabled at the same time. The *enabled* attribute is **false** by default.
12. The *pushStateWhenCoordinator* element is set to **true** by default. If **true**, this property will cause a node that has become the coordinator to transfer in-memory state to the underlying cache store. This parameter is useful where the coordinator has crashed and a new coordinator

is elected.

[Report a bug](#)

#### 17.2.4. About SKIP\_CACHE\_LOAD Flag

In Red Hat JBoss Data Grid's Remote Client-Server mode, when the cache is preloaded from a cache store and eviction is disabled, read requests go to the memory. If the entry is not found in a memory during a read request, it accesses the cache store which may impact the read performance.

To avoid referring to the cache store when a key is not found in the memory, use the **SKIP\_CACHE\_LOAD** flag.

[Report a bug](#)

### 17.3. SHARED CACHE STORES

A shared cache store is a cache store that is shared by multiple cache instances.

A shared cache store is useful when all instances in a cluster communicate with the same remote, shared database using the same JDBC settings. In such an instance, configuring a shared cache store prevents the unnecessary repeated write operations that occur when various cache instances attempt to write the same data to the cache store.

[Report a bug](#)

#### 17.3.1. Invalidation Mode and Shared Cache Stores

When used in conjunction with a shared cache store, Red Hat JBoss Data Grid's invalidation mode causes remote caches to see the shared cache store to retrieve modified data.

The benefits of using invalidation mode in conjunction with shared cache stores include the following:

- Compared to replication messages, which contain the updated data, invalidation messages are much smaller and result in reduced network traffic.
- The remaining cluster caches look up modified data from the shared cache store lazily and only when required to do so, resulting in further reduced network traffic.

[Report a bug](#)

#### 17.3.2. The Cache Store and Cache Passivation

In Red Hat JBoss Data Grid, a cache store can be used to enforce the passivation of entries and to activate eviction in a cache. Whether passivation mode or activation mode are used, the configured cache store both reads from and writes to the data store.

When passivation is disabled in JBoss Data Grid, after the modification, addition or removal of an element is carried out the cache store steps in to persist the changes in the store.

[Report a bug](#)

#### 17.3.3. Application Cachestore Registration



It is not necessary to register an application cache store for an isolated deployment. This is not a requirement in Red Hat JBoss Data Grid because lazy deserialization is used to work around this problem.

[Report a bug](#)

## 17.4. CONNECTION FACTORIES

In Red Hat JBoss Data Grid, all JDBC cache stores rely on a **ConnectionFactory** implementation to obtain a database connection. This process is also known as connection management or pooling.

A connection factory can be specified using the **ConnectionFactoryClass** configuration attribute. JBoss Data Grid includes the following **ConnectionFactory** implementations:

- `ManagedConnectionFactory`
- `SimpleConnectionFactory`.
- `PooledConnectionFactory`.

[Report a bug](#)

### 17.4.1. About `ManagedConnectionFactory`

**ManagedConnectionFactory** is a connection factory that is ideal for use within managed environments such as application servers. This connection factory can explore a configured location in the JNDI tree and delegate connection management to the **DataSource**.

[Report a bug](#)

### 17.4.2. About `SimpleConnectionFactory`

**SimpleConnectionFactory** is a connection factory that creates database connections on a per invocation basis. This connection factory is not designed for use in a production environment.

[Report a bug](#)

## CHAPTER 18. CACHE STORE IMPLEMENTATIONS

The cache store connects Red Hat JBoss Data Grid to the persistent data store. Cache stores are associated with individual caches. Different caches attached to the same cache manager can have different cache store configurations.



### NOTE

If a clustered cache is configured with an unshared cache store (where *shared* is set to **false**), on node join, stale entries which might have been removed from the cluster might still be present in the stores and can reappear.

[Report a bug](#)

### 18.1. CACHE STORE COMPARISON

Select a cache store based on your requirements. The following is a summary of high level differences between the cache stores available in Red Hat JBoss Data Grid:

- The Single File Cache Store is a local file cache store. It persists data locally for each node of the clustered cache. The Single File Cache Store provides superior read and write performance, but keeps keys in memory which limits its use when persisting large data sets at each node. See [Section 18.4, “Single File Cache Store”](#) for details.
- The LevelDB file cache store is a local file cache store which provides high read and write performance. It does not have the limitation of Single File Cache Store of keeping keys in memory. See [Section 18.5, “LevelDB Cache Store”](#) for details.
- The JDBC cache store is a cache store that may be shared, if required. When using it, all nodes of a clustered cache persist to a single database or a local JDBC database for every node in the cluster. The shared cache store lacks the scalability and performance of a local cache store such as the LevelDB cache store, but it provides a single location for persisted data. The JDBC cache store persists entries as binary blobs, which are not readable outside JBoss Data Grid. See [Section 18.6, “JDBC Based Cache Stores”](#) for details.
- The JPA Cache Store (supported in Library mode only) is a shared cache store like JDBC cache store, but preserves schema information when persisting to the database. Therefore, the persisted entries can be read outside JBoss Data Grid. See [Section 18.8, “JPA Cache Store”](#) for details.

[Report a bug](#)

### 18.2. CACHE STORE CONFIGURATION DETAILS (LIBRARY MODE)

The following lists contain details about the configuration elements and parameters for cache store elements in JBoss Data Grid's Library mode:

#### The `namedCache` Element

- Add the name value to the *name* parameter to set the name of the cache store.

#### The `persistence` Element

- The *passivation* parameter affects the way in which Red Hat JBoss Data Grid interacts with

stores. When an object is evicted from in-memory cache, passivation writes it to a secondary data store, such as a system or a database. Valid values for this parameter are **true** and **false** but *passivation* is set to **false** by default.

### The singleFile Element

- The *shared* parameter indicates that the cache store is shared by different cache instances. For example, where all instances in a cluster use the same JDBC settings to talk to the same remote, shared database. *shared* is **false** by default. When set to **true**, it prevents duplicate data being written to the cache store by different cache instances. For the LevelDB cache stores, this parameter must be excluded from the configuration, or set to **false** because sharing this cache store is not supported.
- The *preload* parameter is set to **false** by default. When set to **true** the data stored in the cache store is preloaded into the memory when the cache starts. This allows data in the cache store to be available immediately after startup and avoids cache operations delays as a result of loading data lazily. Preloaded data is only stored locally on the node, and there is no replication or distribution of the preloaded data. Red Hat JBoss Data Grid will only preload up to the maximum configured number of entries in eviction.
- The *fetchPersistentState* parameter determines whether or not to fetch the persistent state of a cache and apply it to the local cache store when joining the cluster. If the cache store is shared the fetch persistent state is ignored, as caches access the same cache store. A configuration exception will be thrown when starting the cache service if more than one cache store has this property set to **true**. The *fetchPersistentState* property is **false** by default.
- The *ignoreModifications* parameter determines whether modification methods are applied to the specific cache store. This allows write operations to be applied to the local file cache store, but not the shared cache store. In some cases, transient application data should only reside in a file-based cache store on the same server as the in-memory cache. For example, this would apply with a further JDBC based cache store used by all servers in the network. *ignoreModifications* is **false** by default.
- The *maxEntries* parameter provides maximum number of entries allowed. The default value is -1 for unlimited entries.
- The *maxKeysInMemory* parameter is used to speed up data lookup. The single file store keeps an index of keys and their positions in the file, restricting the size of the index using the *maxKeysInMemory* parameter. The default value for this parameter is -1.
- The *purgeOnStartup* parameter controls whether cache store is purged when it starts up.
- The *location* configuration element sets a location on disk where the store can write.

### The async Element

The **async** element contains parameters that configure various aspects of the cache store.

- The *enabled* parameter determines whether the file store is asynchronous.
- The *threadPoolSize* parameter specifies the number of threads that concurrently apply modifications to the store. The default value for this parameter is **1**.
- The *flushLockTimeout* parameter specifies the time to acquire the lock which guards the state to be flushed to the cache store periodically. The default value for this parameter is **1**.

- The *modificationQueueSize* parameter specifies the size of the modification queue for the asynchronous store. If updates are made at a rate that is faster than the underlying cache store can process this queue, then the asynchronous store behaves like a synchronous store for that period, blocking until the queue can accept more elements. The default value for this parameter is **1024** elements.
- The *shutdownTimeout* parameter specifies maximum amount of time that can be taken to stop the cache store. Default value for this parameter is **25** seconds.

### The singleton Element

The *singleton* element enables modifications to be stored by only one node in the cluster. This node is called the coordinator. The coordinator pushes the caches in-memory states to disk. The *shared* element cannot be defined with *singleton* enabled at the same time.

- The *enabled* attribute determines whether this feature is enabled. Valid values for this parameter are **true** and **false**. The *enabled* attribute is set to **false** by default.
- The *pushStateWhenCoordinator* parameter is set to **true** by default. If **true**, this property causes a node that has become the coordinator to transfer in-memory state to the underlying cache store. This parameter is useful where the coordinator has crashed and a new coordinator is elected.
- When *pushStateWhenCoordinator* is set to **true**, the *pushStateTimeout* parameter sets the maximum number of milliseconds that the process pushing the in-memory state to the underlying cache loader can take. The default time for this parameter is 10 seconds.

### The remoteStore Element

- The *remoteCacheName* attribute specifies the name of the remote cache to which it intends to connect in the remote Infinispan cluster. The default cache will be used if the remote cache name is unspecified.
- The *fetchPersistentState* attribute, when set to **true**, ensures that the persistent state is fetched when the remote cache joins the cluster. If multiple cache stores are chained, only one cache store can have this property set to **true**. The default for this value is **false**.
- The *shared* attribute is set to **true** when multiple cache instances share a cache store, which prevents multiple cache instances writing the same modification individually. The default for this attribute is **false**.
- The *preload* attribute ensures that the cache store data is pre-loaded into memory and is immediately accessible after starting up. The disadvantage of setting this to **true** is that the start up time increases. The default value for this attribute is **false**.
- The *ignoreModifications* attribute prevents cache modification operations such as put, remove, clear, store, etc. from affecting the cache store. As a result, the cache store can become out of sync with the cache. The default value for this attribute is **false**.
- The *purgeOnStartup* attribute ensures that the cache store is purged during the start up process. The default value for this attribute is **false**.
- The *tcpNoDelay* attribute triggers the **TCP NODELAY** stack. The default value for this attribute is **true**.

- The ***pingOnStartup*** attribute sends a ping request to a back end server to fetch the cluster topology. The default value for this attribute is **true**.
- The ***keySizeEstimate*** attribute provides an estimation of the key size. The default value for this attribute is **64**.
- The ***valueSizeEstimate*** attribute specifies the size of the byte buffers when serializing and deserializing values. The default value for this attribute is **512**.
- The ***forceReturnValues*** attribute sets whether **FORCE\_RETURN\_VALUE** is enabled for all calls. The default value for this attribute is **false**.

### The servers and server Elements

Create a **servers** element within the **remoteStore** element to set up the server information for multiple servers. Add a **server** element within the general **servers** element to add the information for a single server.

- The ***host*** attribute configures the host address.
- The ***port*** attribute configures the port used by the Remote Cache Store.

### The connectionPool Element

- The ***maxActive*** attribute indicates the maximum number of active connections for each server at a time. The default value for this attribute is **-1** which indicates an infinite number of active connections.
- The ***maxIdle*** attribute indicates the maximum number of idle connections for each server at a time. The default value for this attribute is **-1** which indicates an infinite number of idle connections.
- The ***maxTotal*** attribute indicates the maximum number of persistent connections within the combined set of servers. The default setting for this attribute is **-1** which indicates an infinite number of connections.
- The ***connectionUrl*** parameter specifies the JDBC driver-specific connection URL.
- The ***username*** parameter contains the username used to connect via the ***connectionUrl***.
- The ***driverClass*** parameter specifies the class name of the driver used to connect to the database.

### The leveldbStore Element

- The ***location*** parameter specifies the location to store the primary cache store. The directory is automatically created if it does not exist.
- The ***expiredLocation*** parameter specifies the location for expired data using. The directory stores expired data before it is purged. The directory is automatically created if it does not exist.
- The ***shared*** parameter specifies whether the cache store is shared. The only supported value for this parameter in the LevelDB cache store is **false**.
- The ***preload*** parameter specifies whether the cache store will be pre-loaded. Valid values are **true** and **false**.

## The `jpaStore` Element

- The ***persistenceUnitName*** attribute specifies the name of the JPA cache store.
- The ***entityClassName*** attribute specifies the fully qualified class name of the JPA entity used to store the cache entry value.
- The ***batchSize*** (optional) attribute specifies the batch size for cache store streaming. The default value for this attribute is **100**.
- The ***storeMetadata*** (optional) attribute specifies whether the cache store keeps the metadata (for example expiration and versioning information) with the entries. The default value for this attribute is **true**.

## The `binaryKeyedJdbcStore`, `stringKeyedJdbcStore`, and `mixedKeyedJdbcStore` Elements

- The ***fetchPersistentState*** parameter determines whether the persistent state is fetched when joining a cluster. Set this to **true** if using a replication and invalidation in a clustered environment. Additionally, if multiple cache stores are chained, only one cache store can have this property enabled. If a shared cache store is used, the cache does not allow a persistent state transfer despite this property being set to **true**. The ***fetchPersistentState*** parameter is **false** by default.
- The ***ignoreModifications*** parameter determines whether operations that modify the cache (e.g. put, remove, clear, store, etc.) do not affect the cache store. As a result, the cache store can become out of sync with the cache.
- The ***purgeOnStartup*** parameter specifies whether the cache store is purged when initially started.
- The ***key2StringMapper*** parameter specifies the class name of the `Key2StringMapper` used to map keys to strings for the database tables.

## The `binaryKeyedTable` and `stringKeyedTable` Elements

- The ***dropOnExit*** parameter specifies whether the database tables are dropped upon shutdown.
- The ***createOnStart*** parameter specifies whether the database tables are created by the store on startup.
- The ***prefix*** parameter defines the string prepended to name of the target cache when composing the name of the cache bucket table.

## The `idColumn`, `dataColumn`, and `timestampColumn` Elements

- The ***name*** parameter specifies the name of the column used.
- The ***type*** parameter specifies the type of the column used.

## The `store` Element

- The ***class*** parameter specifies the class name of the cache store implementation.
- The ***preload*** parameter specifies whether to load entries into the cache during start up. Valid values for this parameter are **true** and **false**.

- The **shared** parameter specifies whether the cache store is shared. This is used when multiple cache instances share a cache store. Valid values for this parameter are **true** and **false**.

### The property Element

- The **name** parameter specifies the name of the property.
- The **value** parameter specifies the value of the property.

[Report a bug](#)

## 18.3. CACHE STORE CONFIGURATION DETAILS (REMOTE CLIENT-SERVER MODE)

The following tables contain details about the configuration elements and parameters for cache store elements in JBoss Data Grid's Remote Client-Server mode:

### The local-cache Element

- The **name** parameter of the **local-cache** attribute is used to specify a name for the cache.
- The **statistics** parameter specifies whether statistics are enabled at the container level. Enable or disable statistics on a per-cache basis by setting the **statistics** attribute to **false**.

### The file-store Element

- The **name** parameter of the **file-store** element is used to specify a name for the file store.
- The **passivation** parameter determines whether entries in the cache are passivated (**true**) or if the cache store retains a copy of the contents in memory (**false**).
- The **purge** parameter specifies whether or not the cache store is purged when it is started. Valid values for this parameter are **true** and **false**.
- The **shared** parameter is used when multiple cache instances share a cache store. This parameter can be set to prevent multiple cache instances writing the same modification multiple times. Valid values for this parameter are **true** and **false**. However, the **shared** parameter is not recommended for the LevelDB cache store because this cache store cannot be shared.
- The **relative-to** property is the directory where the **file-store** stores the data. It is used to define a named path.
- The **path** property is the name of the file where the data is stored. It is a relative path name that is appended to the value of the **relative-to** property to determine the complete path.
- The **maxEntries** parameter provides maximum number of entries allowed. The default value is -1 for unlimited entries.
- The **fetch-state** parameter when set to true fetches the persistent state when joining a cluster. If multiple cache stores are chained, only one of them can have this property enabled. Persistent state transfer with a shared cache store does not make sense, as the same persistent store that provides the data will just end up receiving it. Therefore, if a shared cache store is used, the cache does not allow a persistent state transfer even if a cache store has this property set to **true**. It is recommended to set this property to true only in a clustered environment. The default value for this parameter is false.

- The ***preload*** parameter when set to true, loads the data stored in the cache store into memory when the cache starts. However, setting this parameter to true affects the performance as the startup time is increased. The default value for this parameter is false.
- The ***singleton*** parameter enables a singleton store cache store. SingletonStore is a delegating cache store used when only one instance in a cluster can interact with the underlying store. However, ***singleton*** parameter is not recommended for **file-store**.

### The store Element

- The ***class*** parameter specifies the class name of the cache store implementation.

### The property Element

- The ***name*** parameter specifies the name of the property.
- The ***value*** parameter specifies the value assigned to the property.

### The remote-store Element

- The ***cache*** parameter defines the name for the remote cache. If left undefined, the default cache is used instead.
- The ***socket-timeout*** parameter sets whether the value defined in ***SO\_TIMEOUT*** (in milliseconds) applies to remote Hot Rod servers on the specified timeout. A timeout value of **0** indicates an infinite timeout.
- The ***tcp-no-delay*** sets whether **TCP\_NODELAY** applies on socket connections to remote Hot Rod servers.
- The ***hotrod-wrapping*** sets whether a wrapper is required for Hot Rod on the remote store.

### The remote-server Element

- The ***outbound-socket-binding*** parameter sets the outbound socket binding for the remote server.

### The binary-keyed-jdbc-store, string-keyed-jdbc-store, and mixed-keyed-jdbc-store Elements

- The ***datasource*** parameter defines the name of a JNDI for the datasource.
- The ***passivation*** parameter determines whether entries in the cache are passivated (**true**) or if the cache store retains a copy of the contents in memory (**false**).
- The ***preload*** parameter specifies whether to load entries into the cache during start up. Valid values for this parameter are **true** and **false**.
- The ***purge*** parameter specifies whether or not the cache store is purged when it is started. Valid values for this parameter are **true** and **false**.
- The ***shared*** parameter is used when multiple cache instances share a cache store. This parameter can be set to prevent multiple cache instances writing the same modification multiple times. Valid values for this parameter are **true** and **false**.



- The *singleton* parameter enables a singleton store cache store. SingletonStore is a delegating cache store used when only one instance in a cluster can interact with the underlying store

### The binary-keyed-table and string-keyed-table Elements

- The *prefix* parameter specifies a prefix string for the database table name.

### The id-column, data-column, and timestamp-column Elements

- The *name* parameter specifies the name of the database column.
- The *type* parameter specifies the type of the database column.

### The leveldb-store Element

- The *path* parameter The directory within the path specified in the *relative-to* parameter where the cache state is stored. If undefined, the path defaults to the cache container name.
- The *passivation* parameter specifies whether passivation is enabled for the LevelDB cache store. Valid values are **true** and **false**.
- The *purge* parameter specifies whether the cache store is purged when it starts up. Valid values are **true** and **false**.

[Report a bug](#)

## 18.4. SINGLE FILE CACHE STORE

Red Hat JBoss Data Grid includes one file system based cache store: the **SingleFileCacheStore**.

The **SingleFileCacheStore** is a simple, file system based implementation and a replacement to the older file system based cache store: the **FileCacheStore**.

**SingleFileCacheStore** stores all key/value pairs and their corresponding metadata information in a single file. To speed up data location, it also keeps all keys and the positions of their values and metadata in memory. Hence, using the single file cache store slightly increases the memory required, depending on the key size and the amount of keys stored. Hence **SingleFileCacheStore** is not recommended for use cases where the keys are too big.

To reduce memory consumption, the size of the cache store can be set to a fixed number of entries to store in the file. However, this works only when Infinispan is used as a cache. When Infinispan used this way, data which is not present in Infinispan can be recomputed or re-retrieved from the authoritative data store and stored in Infinispan cache. The reason for this limitation is because once the maximum number of entries is reached, older data in the cache store is removed, so if Infinispan was used as an authoritative data store, it would lead to data loss which is undesirable in this use case

Due to its limitations, **SingleFileCacheStore** can be used in a limited capacity in production environments. It can not be used on shared file system (such as NFS and Windows shares) due to a lack of proper file locking, resulting in data corruption. Furthermore, file systems are not inherently transactional, resulting in file writing failures during the commit phase if the cache is used in a transactional context.

[Report a bug](#)

### 18.4.1. Single File Store Configuration (Remote Client-Server Mode)

The following is an example of a Single File Store configuration for Red Hat JBoss Data Grid's Remote Client-Server mode:

```
<local-cache name="default" statistics="true">
  <file-store name="myFileStore"
    passivation="true"
    purge="true"
    relative-to="{PATH}"
    path="{DIRECTORY}"
    max-entries="10000"
    fetch-state="true"
    preload="false" />
</local-cache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, "Cache Store Configuration Details \(Remote Client-Server Mode\)"](#).

[Report a bug](#)

### 18.4.2. Single File Store Configuration (Library Mode)

In Red Hat JBoss Grid's Library mode, configure a Single File Cache Store as follows:.

```
<namedCache name="writeThroughToFile">
  <persistence passivation="false">
    <singleFile fetchPersistentState="true"
      ignoreModifications="false"
      purgeOnStartup="false"
      shared="false"
      preload="false"
      location="/tmp/Another-FileCacheStore-Location"
      maxEntries="100"
      maxKeysInMemory="100">
      <async enabled="true"
        threadPoolSize="500"
        flushLockTimeout="1"
        modificationQueueSize="1024"
        shutdownTimeout="25000"/>
    </singleFile>
  </persistence>
</namedCache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, "Cache Store Configuration Details \(Library Mode\)"](#).

[Report a bug](#)

### 18.4.3. Migrating data from FileCacheStore to SingleFileCacheStore

Red Hat JBoss Data Grid stores data in a different format than previous versions of JBoss Data Grid. As a result, the newer version of JBoss Data Grid cannot read data stored by older versions. Use rolling upgrades to upgrade persisted data from the format used by the old JBoss Data Grid to the new format.

Additionally, the newer version of JBoss Data Grid also stores persistence configuration information in a different location.

Rolling upgrades is the process by which a JBoss Data Grid installation is upgraded without a service shutdown. In Library mode, it refers to a node installation where JBoss Data Grid is running in Library mode. For JBoss Data Grid servers, it refers to the server side components. The upgrade can be due to either hardware or software change such as upgrading JBoss Data Grid.

Rolling upgrades are only available in JBoss Data Grid's Remote Client-Server mode.

[Report a bug](#)

## 18.5. LEVELDB CACHE STORE

LevelDB is a key-value storage engine that provides an ordered mapping from string keys to string values.

The LevelDB Cache Store uses two filesystem directories. Each directory is configured for a LevelDB database. One directory stores the non-expired data and the second directory stores the expired keys before a purge.

[Report a bug](#)

### 18.5.1. Configuring LevelDB Cache Store (Remote Client-Server Mode)

**Procedure 18.1. To configure LevelDB Cache Store:**

- Add the following elements to a cache definition in `standalone.xml` to configure the database:

```
<leveldb-store path="/path/to/leveldb/data"
              passivation="false"
              purge="false" >
  <expiration path="/path/to/leveldb/expires/data" />
  <implementation type="JNI" />
</leveldb-store>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, "Cache Store Configuration Details \(Remote Client-Server Mode\)"](#).

[Report a bug](#)

### 18.5.2. LevelDB Cache Store Programmatic Configuration

The following is a sample programmatic configuration of LevelDB Cache Store:

```
Configuration cacheConfig = new ConfigurationBuilder().persistence()
    .addStore(LevelDBStoreConfigurationBuilder.class)
    .location("/tmp/leveldb/data")
    .expiredLocation("/tmp/leveldb/expired").build();
```

**Procedure 18.2. LevelDB Cache Store programmatic configuration**

1. Use the *ConfigurationBuilder* to create a new configuration object.

2. Add the store using `LevelDBCacheStoreConfigurationBuilder` class to build its configuration.
3. Set the LevelDB Cache Store location path. The specified path stores the primary cache store data. The directory is automatically created if it does not exist.
4. Specify the location for expired data using the `expiredLocation` parameter for the LevelDB Store. The specified path stores expired data before it is purged. The directory is automatically created if it does not exist.

**NOTE**

Programmatic configurations can only be used with Red Hat JBoss Data Grid Library mode.

[Report a bug](#)

### 18.5.3. LevelDB Cache Store Sample XML Configuration (Library Mode)

The following is a sample XML configuration of LevelDB Cache Store:

```
<namedCache name="vehicleCache">
  <persistence passivation="false">
    <levelDbStore xmlns="urn:infinispan:config:store:leveldb:6.0"
      location="/path/to/leveldb/data"
      expiredLocation="/path/to/expired/data"
      shared="false"
      preload="true"/>
  </persistence>
</namedCache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).

[Report a bug](#)

### 18.5.4. Configure a LevelDB Cache Store Using JBoss Operations Network

Use the following procedure to set up a new LevelDB cache store using the JBoss Operations Network.

**Procedure 18.3.**

1. Ensure that Red Hat JBoss Operations Network 3.2 or better is installed and started.
2. Install the Red Hat JBoss Data Grid Plugin Pack for JBoss Operations Network 3.2.0.
3. Ensure that JBoss Data Grid is installed and started.
4. Import JBoss Data Grid server into the inventory.
5. Configure the JBoss Data Grid connection settings.
6. Create a new LevelDB cache store as follows:

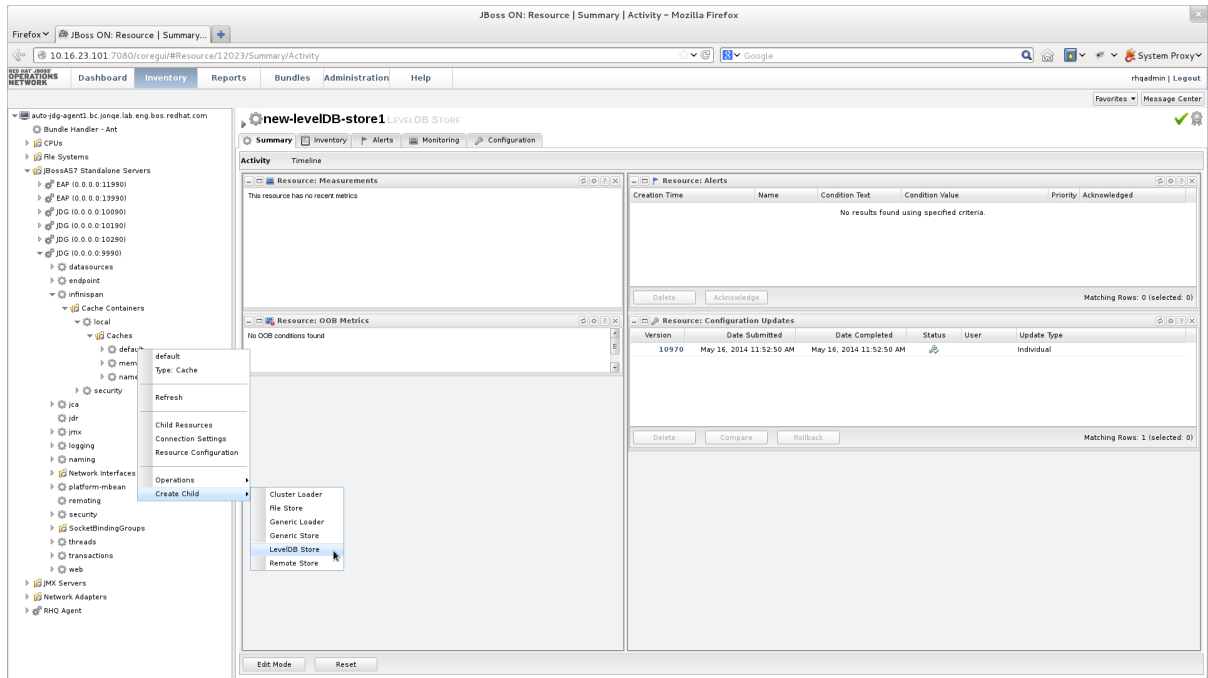


Figure 18.1. Create a new LevelDB Cache Store

- a. Right-click the **default** cache.
- b. In the menu, mouse over the **Create Child** option.
- c. In the submenu, click **LevelDB Store**.

7. Name the new LevelDB cache store as follows:

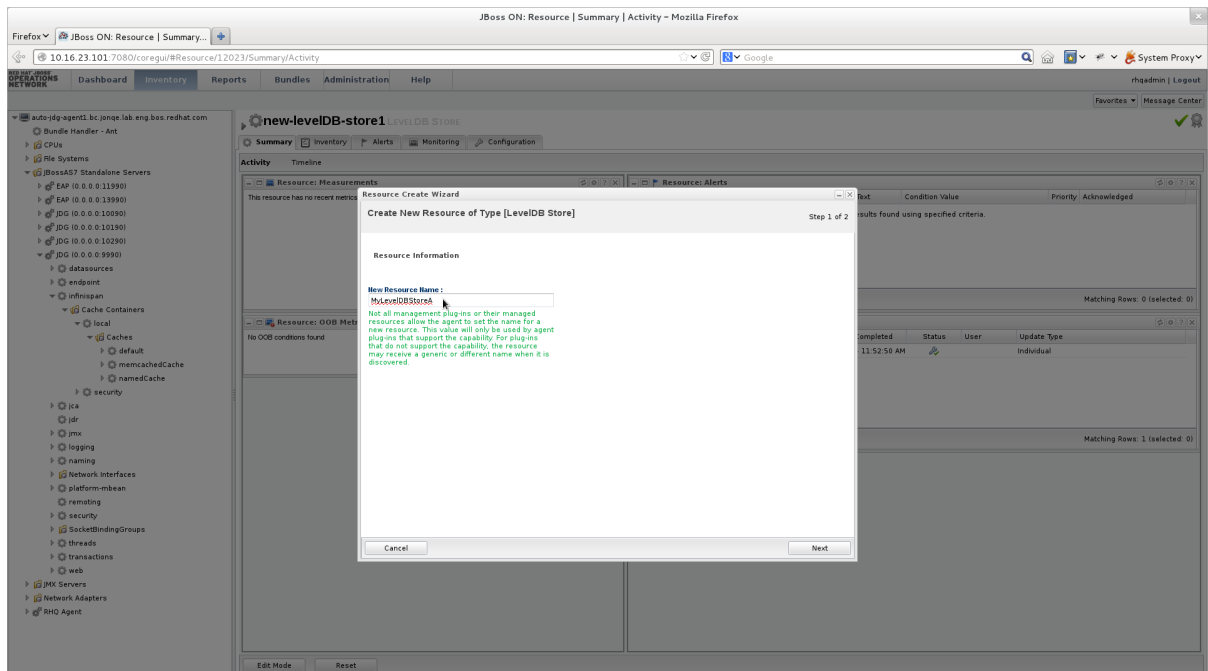


Figure 18.2. Name the new LevelDB Cache Store

- a. In the **Resource Create Wizard** that appears, add a name for the new LevelDB Cache Store.
- b. Click **Next** to continue.

## 8. Configure the LevelDB Cache Store settings as follows:

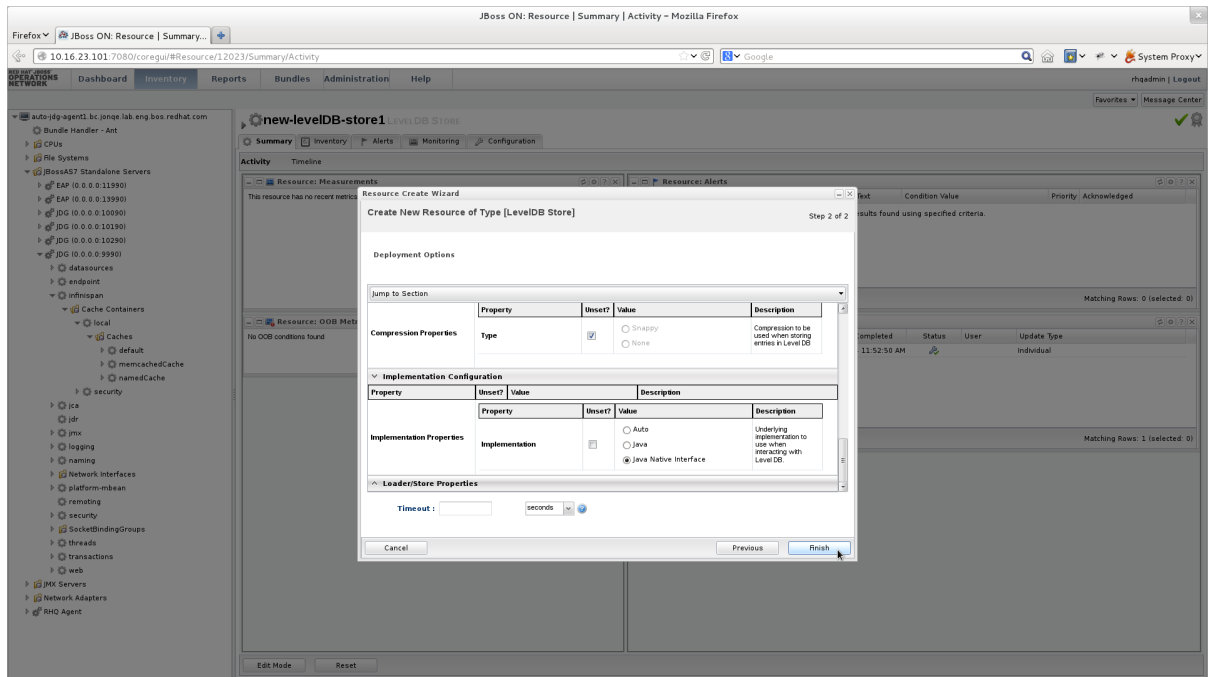


Figure 18.3. Configure the LevelDB Cache Store Settings

- Use the options in the configuration window to configure a new LevelDB cache store.
- Click **Finish** to complete the configuration.

## 9. Schedule a restart operation as follows:

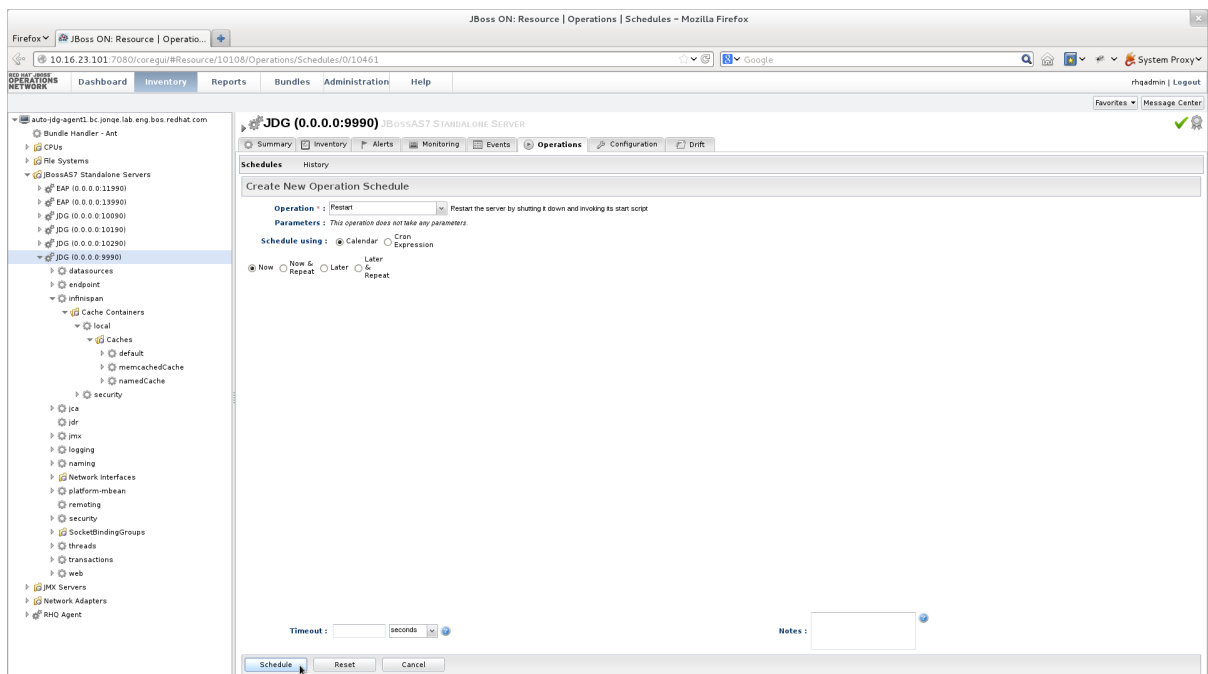



Figure 18.4. Schedule a Restart Operation

- In the screen's left panel, expand the **JBossAS7 Standalone Servers** entry, if it is not currently expanded.
- Click **JDG (0.0.0.0:9990)** from the expanded menu items.

- c. In the screen's right panel, details about the selected server display. Click the **Operations** tab.
  - d. In the **Operation** drop-down box, select the **Restart** operation.
  - e. Select the radio button for the **Now** entry.
  - f. Click **Schedule** to restart the server immediately.
10. Discover the new LevelDB cache store as follows:



**Figure 18.5. Discover the New LevelDB Cache Store**

- a. In the screen's left panel, select each of the following items in the specified order to expand them: **JBossAS7 Standalone Servers** → **JDG (0.0.0.0:9990)** → **infinispan** → **Cache Containers** → **local** → **Caches** → **default** → **LevelDB Stores**
- b. Click the name of your new LevelDB Cache Store to view its configuration information in the right panel.

[Report a bug](#)

## 18.6. JDBC BASED CACHE STORES

Red Hat JBoss Data Grid offers several cache stores for use with common data storage formats. JDBC based cache stores are used with any cache store that exposes a JDBC driver. JBoss Data Grid offers the following JDBC based cache stores depending on the key to be persisted:

- **JdbcBinaryStore.**
- **JdbcStringBasedStore.**
- **JdbcMixedStore.**

[Report a bug](#)

## 18.6.1. JdbcBinaryStores

The **JdbcBinaryStore** supports all key types. It stores all keys with the same hash value (**hashCode** method on the key) in the same table row/blob. The hash value common to the included keys is set as the primary key for the table row/blob. As a result of this hash value, **JdbcBinaryStore** offers excellent flexibility but at the cost of concurrency and throughput.

As an example, if three keys (**k1**, **k2** and **k3**) have the same hash code, they are stored in the same table row. If three different threads attempt to concurrently update **k1**, **k2** and **k3**, they must do it sequentially because all three keys share the same row and therefore cannot be simultaneously updated.

[Report a bug](#)

### 18.6.1.1. JdbcBinaryStore Configuration (Remote Client-Server Mode)

The following is a configuration for **JdbcBinaryStore** using Red Hat JBoss Data Grid's Remote Client-Server mode with Passivation enabled:

```
<local-cache name="customCache">

  <!-- Additional configuration elements here -->
  <binary-keyed-jdbc-store datasource="java:jboss/datasources/JdbcDS"

    passivation="${true/false}"
    preload="${true/false}"
    purge="${true/false}">
    <binary-keyed-table prefix="JDG">
      <id-column name="id"
        type="${id.column.type}"/>
      <data-column name="datum"
        type="${data.column.type}"/>
      <timestamp-column name="version"
        type="${timestamp.column.type}"/>
    </binary-keyed-table>
  </binary-keyed-jdbc-store>
</local-cache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, "Cache Store Configuration Details \(Remote Client-Server Mode\)"](#).

[Report a bug](#)

### 18.6.1.2. JdbcBinaryStore Configuration (Library Mode)

The following is a sample configuration for the **JdbcBinaryStore**:

```
<infinispan
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:infinispan:config:6.0
http://www.infinispan.org/schemas/infinispan-config-6.0.xsd
  urn:infinispan:config:jdbc:6.0
http://www.infinispan.org/schemas/infinispan-cachestore-jdbc-config-
6.0.xsd"
  xmlns="urn:infinispan:config:6.0">
```



```

        <!-- Additional configuration elements here -->
<persistence>
<binaryKeyedJdbcStore xmlns="urn:infinispan:config:jdbc:6.0"
                        fetchPersistentState="false"
                        ignoreModifications="false"
                        purgeOnStartup="false">
  <connectionPool
connectionUrl="jdbc:h2:mem:infinispan_binary_based;DB_CLOSE_DELAY=-1"
  username="sa"
  driverClass="org.h2.Driver"/>
  <binaryKeyedTable dropOnExit="true"
    createOnStart="true"
    prefix="ISPN_BUCKET_TABLE">
    <idColumn name="ID_COLUMN"
      type="VARCHAR(255)" />
    <dataColumn name="DATA_COLUMN"
      type="BINARY" />
    <timestampColumn name="TIMESTAMP_COLUMN"
      type="BIGINT" />
  </binaryKeyedTable>
</binaryKeyedJdbcStore>
</persistence>

```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).

[Report a bug](#)

### 18.6.1.3. JdbcBinaryStore Programmatic Configuration

The following is a sample configuration for the `JdbcBinaryStore`:

```

ConfigurationBuilder builder = new ConfigurationBuilder();
builder.persistence()
  .addStore(JdbcBinaryStoreConfigurationBuilder.class)
  .fetchPersistentState(false)
  .ignoreModifications(false)
  .purgeOnStartup(false)
  .table()
    .dropOnExit(true)
    .createOnStart(true)
    .tableNamePrefix("ISPN_BUCKET_TABLE")
    .idColumnName("ID_COLUMN").idColumnType("VARCHAR(255)")
    .dataColumnName("DATA_COLUMN").dataColumnType("BINARY")

  .timestampColumnName("TIMESTAMP_COLUMN").timestampColumnType("BIGINT")
  .connectionPool()

  .connectionUrl("jdbc:h2:mem:infinispan_binary_based;DB_CLOSE_DELAY=-1")
    .username("sa")
    .driverClass("org.h2.Driver");

```

#### Procedure 18.4. JdbcBinaryStore Programmatic Configuration (Library Mode)

1. Use the *ConfigurationBuilder* to create a new configuration object.

2. Add the **JdbcBinaryStore** configuration builder to build a specific configuration related to this store.
3. The **fetchPersistentState** element determines whether or not to fetch the persistent state of a cache and apply it to the local cache store when joining the cluster. If the cache store is shared the fetch persistent state is ignored, as caches access the same cache store. A configuration exception will be thrown when starting the cache service if more than one cache loader has this property set to **true**. The **fetchPersistentState** property is **false** by default.
4. The **ignoreModifications** element determines whether write methods are pushed to the specific cache loader by allowing write operations to the local file cache loader, but not the shared cache loader. In some cases, transient application data should only reside in a file-based cache loader on the same server as the in-memory cache. For example, this would apply with a further JDBC based cache loader used by all servers in the network. **ignoreModifications** is **false** by default.
5. The **purgeOnStartup** element specifies whether the cache is purged when initially started.
6. Configure the table as follows:
  - a. **dropOnExit** determines if the table will be dropped when the cache store is stopped. This is set to **false** by default.
  - b. **createOnStart** creates the table when starting the cache store if no table currently exists. This method is **true** by default.
  - c. **tableNamePrefix** sets the prefix for the name of the table in which the data will be stored.
  - d. The **idColumnName** property defines the column where the cache key or bucket ID is stored.
  - e. The **dataColumnName** property specifies the column where the cache entry or bucket is stored.
  - f. The **timestampColumnName** element specifies the column where the time stamp of the cache entry or bucket is stored.
7. The **connectionPool** element specifies a connection pool for the JDBC driver using the following parameters:
  - a. The **connectionUrl** parameter specifies the JDBC driver-specific connection URL.
  - b. The **username** parameter contains the user name used to connect via the **connectionUrl**.
  - c. The **driverClass** parameter specifies the class name of the driver used to connect to the database.

**NOTE**

Programmatic configurations can only be used with Red Hat JBoss Data Grid Library mode.

## 18.6.2. JdbcStringBasedStores

The **JdbcStringBasedStore** stores each entry in its own row in the table, instead of grouping multiple entries into each row, resulting in increased throughput under a concurrent load. It also uses a (pluggable) bijection that maps each key to a **String** object. The **Key2StringMapper** interface defines the bijection.

Red Hat JBoss Data Grid includes a default implementation called **DefaultTwoWayKey2StringMapper** that handles primitive types.

[Report a bug](#)

### 18.6.2.1. JdbcStringBasedStore Configuration (Remote Client-Server Mode)

The following is a sample **JdbcStringBasedStore** for Red Hat JBoss Data Grid's Remote Client-Server mode:

```
<local-cache name="customCache">
  <!-- Additional configuration elements here -->
  <string-keyed-jdbc-store datasource="java:jboss/datasources/JdbcDS"
    passivation="true"
    preload="false"
    purge="false"
    shared="false"
    singleton="true">
    <string-keyed-table prefix="JDG">
      <id-column name="id"
        type="{id.column.type}"/>
      <data-column name="datum"
        type="{data.column.type}"/>
      <timestamp-column name="version"
        type="{timestamp.column.type}"/>
    </string-keyed-table>
  </string-keyed-jdbc-store>
</local-cache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, "Cache Store Configuration Details \(Remote Client-Server Mode\)"](#).

[Report a bug](#)

### 18.6.2.2. JdbcStringBasedStore Configuration (Library Mode)

The following is a sample configuration for the **JdbcStringBasedStore**:

```
<infinispan
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:infinispan:config:6.0
    http://www.infinispan.org/schemas/infinispan-config-6.0.xsd
    urn:infinispan:config:jdbc:6.0
    http://www.infinispan.org/schemas/infinispan-cachestore-jdbc-config-
    6.0.xsd"
  xmlns="urn:infinispan:config:6.0">
  <!-- Additional configuration elements here -->
  <persistence>
```

```

<stringKeyedJdbcStore xmlns="urn:infinispan:config:jdbc:6.0"
                      fetchPersistentState="false"
                      ignoreModifications="false"
                      purgeOnStartup="false"

key2StringMapper="org.infinispan.loaders.keymappers.DefaultTwoWayKey2StringMapper">
  <connectionPool
connectionUrl="jdbc:h2:mem:infinispan_binary_based;DB_CLOSE_DELAY=-1"
  username="sa"
  driverClass="org.h2.Driver"/>
  <stringKeyedTable dropOnExit="true"
  createOnStart="true"
  prefix="ISPN_STRING_TABLE">
    <idColumn name="ID_COLUMN"
  type="VARCHAR(255)" />
    <dataColumn name="DATA_COLUMN"
  type="BINARY" />
    <timestampColumn name="TIMESTAMP_COLUMN"
  type="BIGINT" />
  </stringKeyedTable>
</stringKeyedJdbcStore>
</persistence>

```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).

[Report a bug](#)

### 18.6.2.3. JdbcStringBasedStore Multiple Node Configuration (Remote Client-Server Mode)

The following is a configuration for the **JdbcStringBasedStore** in Red Hat JBoss Data Grid's Remote Client-Server mode. This configuration is used when multiple nodes must be used.

```

<subsystem xmlns="urn:infinispan:server:core:6.1" default-cache-
container="default">
  <cache-container <!-- Additional configuration information here --> >
    <!-- Additional configuration elements here -->
    <replicated-cache>
      <!-- Additional configuration elements here -->
      <string-keyed-jdbc-store datasource="java:jboss/datasources/JdbcDS"
        fetch-state="true"
        passivation="false"
        preload="false"
        purge="false"
        shared="false"
        singleton="true">
        <string-keyed-table prefix="JDG">
          <id-column name="id"
            type="{id.column.type}"/>
          <data-column name="datum"
            type="{data.column.type}"/>
          <timestamp-column name="version"
            type="{timestamp.column.type}"/>
        </string-keyed-table>

```

```

    </string-keyed-jdbc-store>
  </replicated-cache>
</cache-container>
</subsystem>

```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, “Cache Store Configuration Details \(Remote Client-Server Mode\)”](#).

[Report a bug](#)

#### 18.6.2.4. JdbcStringBasedStore Programmatic Configuration

The following is a sample configuration for the `JdbcStringBasedStore`:

```

ConfigurationBuilder builder = new ConfigurationBuilder();

builder.persistence().addStore(JdbcStringBasedStoreConfigurationBuilder.class)
    .fetchPersistentState(false)
    .ignoreModifications(false)
    .purgeOnStartup(false)
    .table()
        .dropOnExit(true)
        .createOnStart(true)
        .tableNamePrefix("ISPN_STRING_TABLE")
        .idColumnName("ID_COLUMN").idColumnType("VARCHAR(255)")
        .dataColumnName("DATA_COLUMN").dataColumnType("BINARY")

    .timestampColumnName("TIMESTAMP_COLUMN").timestampColumnType("BIGINT")
    .connectionPool()


    .connectionUrl("jdbc:h2:mem:infinispan_binary_based;DB_CLOSE_DELAY=-1")
        .username("sa")
        .driverClass("org.h2.Driver");

```

#### Procedure 18.5. Configure the JdbcStringBasedStore Programmatically

1. Use the *ConfigurationBuilder* to create a new configuration object.
2. Add the `JdbcStringBasedStore` configuration builder to build a specific configuration related to this store.
3. The *fetchPersistentState* parameter determines whether or not to fetch the persistent state of a cache and apply it to the local cache store when joining the cluster. If the cache store is shared the fetch persistent state is ignored, as caches access the same cache store. A configuration exception will be thrown when starting the cache service if more than one cache loader has this property set to `true`. The *fetchPersistentState* property is `false` by default.
4. The *ignoreModifications* parameter determines whether write methods are pushed to the specific cache loader by allowing write operations to the local file cache loader, but not the shared cache loader. In some cases, transient application data should only reside in a file-based cache loader on the same server as the in-memory cache. For example, this would apply with a further JDBC based cache loader used by all servers in the network. *ignoreModifications* is `false` by default.

5. The ***purgeOnStartup*** parameter specifies whether the cache is purged when initially started.
6. Configure the Table
  - a. ***dropOnExit*** determines if the table will be dropped when the cache store is stopped. This is set to ***false*** by default.
  - b. ***createOnStart*** creates the table when starting the cache store if no table currently exists. This method is ***true*** by default.
  - c. ***tableNamePrefix*** sets the prefix for the name of the table in which the data will be stored.
  - d. The ***idColumnName*** property defines the column where the cache key or bucket ID is stored.
  - e. The ***dataColumnName*** property specifies the column where the cache entry or bucket is stored.
  - f. The ***timestampColumnName*** element specifies the column where the time stamp of the cache entry or bucket is stored.
7. The ***connectionPool*** element specifies a connection pool for the JDBC driver using the following parameters:
  - a. The ***connectionUrl*** parameter specifies the JDBC driver-specific connection URL.
  - b. The ***username*** parameter contains the username used to connect via the ***connectionUrl***.
  - c. The ***driverClass*** parameter specifies the class name of the driver used to connect to the database.

**NOTE**

Programmatic configurations can only be used with Red Hat JBoss Data Grid Library mode.

[Report a bug](#)

### 18.6.3. JdbcMixedStores

The **JdbcMixedStore** is a hybrid implementation that delegates keys based on their type to either the **JdbcBinaryStore** or **JdbcStringBasedStore**.

[Report a bug](#)

#### 18.6.3.1. JdbcMixedStore Configuration (Remote Client-Server Mode)

The following is a configuration for a **JdbcMixedStore** for Red Hat JBoss Data Grid's Remote Client-Server mode:

```
<local-cache name="customCache">  
  <mixed-keyed-jdbc-store datasource="java:jboss/datasources/JdbcDS"  
    passivation="true"
```

```

    preload="false"
    purge="false">
<binary-keyed-table prefix="MIX_BKT2">
  <id-column name="id"
    type="{id.column.type}"/>
  <data-column name="datum"
    type="{data.column.type}"/>
  <timestamp-column name="version"
    type="{timestamp.column.type}"/>
</binary-keyed-table>
<string-keyed-table prefix="MIX_STR2">
  <id-column name="id"
    type="{id.column.type}"/>
  <data-column name="datum"
    type="{data.column.type}"/>
  <timestamp-column name="version"
    type="{timestamp.column.type}"/>
</string-keyed-table>
</mixed-keyed-jdbc-store>
</local-cache>

```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, “Cache Store Configuration Details \(Remote Client-Server Mode\)”](#).

[Report a bug](#)

### 18.6.3.2. JdbcMixedStore Configuration (Library Mode)

The following is a sample configuration for the `mixedKeyedJdbcStore`:

```

<infinispan
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="urn:infinispan:config:6.0
http://www.infinispan.org/schemas/infinispan-config-6.0.xsd
  urn:infinispan:config:jdbc:6.0
http://www.infinispan.org/schemas/infinispan-cachestore-jdbc-config-
6.0.xsd"
  xmlns="urn:infinispan:config:6.0">
  <!-- Additional configuration elements here -->
  <persistence>
  <mixedKeyedJdbcStore xmlns="urn:infinispan:config:jdbc:6.0"
    fetchPersistentState="false"
    ignoreModifications="false"
    purgeOnStartup="false"

key2StringMapper="org.infinispan.persistence.keymappers.DefaultTwoWayKey2S
tringMapper">
  <connectionPool
connectionUrl="jdbc:h2:mem:infinispan_binary_based;DB_CLOSE_DELAY=-1"
  username="sa"
  driverClass="org.h2.Driver"/>
  <binaryKeyedTable dropOnExit="true"
    createOnStart="true"
    prefix="ISPN_BUCKET_TABLE_BINARY">
  <idColumn name="ID_COLUMN"
    type="VARCHAR(255)" />

```

```

<dataColumn name="DATA_COLUMN"
  type="BINARY" />
<timestampColumn name="TIMESTAMP_COLUMN"
  type="BIGINT" />
</binaryKeyedTable>
<stringKeyedTable dropOnExit="true"
  createOnStart="true"
  prefix="ISPN_BUCKET_TABLE_STRING">
  <idColumn name="ID_COLUMN"
    type="VARCHAR(255)" />
  <dataColumn name="DATA_COLUMN"
    type="BINARY" />
  <timestampColumn name="TIMESTAMP_COLUMN"
    type="BIGINT" />
</stringKeyedTable>
</mixedKeyedJdbcStore>
</persistence>

```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).

[Report a bug](#)

### 18.6.3.3. JdbcMixedStore Programmatic Configuration

The following is a sample configuration for the `JdbcMixedStore`:

```

ConfigurationBuilder builder = new ConfigurationBuilder();
builder.persistence().addStore(JdbcMixedStoreConfigurationBuilder.class)
  .fetchPersistentState(false)
  .ignoreModifications(false)
  .purgeOnStartup(false)
  .stringTable()
    .dropOnExit(true)
    .createOnStart(true)
    .tableNamePrefix("ISPN_MIXED_STR_TABLE")
    .idColumnName("ID_COLUMN").idColumnType("VARCHAR(255)")
    .dataColumnName("DATA_COLUMN").dataColumnType("BINARY")

  .timestampColumnName("TIMESTAMP_COLUMN").timestampColumnType("BIGINT")
    .binaryTable()
      .dropOnExit(true)
      .createOnStart(true)
      .tableNamePrefix("ISPN_MIXED_BINARY_TABLE")
      .idColumnName("ID_COLUMN").idColumnType("VARCHAR(255)")
      .dataColumnName("DATA_COLUMN").dataColumnType("BINARY")

  .timestampColumnName("TIMESTAMP_COLUMN").timestampColumnType("BIGINT")
    .connectionPool()

  .connectionUrl("jdbc:h2:mem:infinispan_binary_based;DB_CLOSE_DELAY=-1")
    .username("sa")
    .driverClass("org.h2.Driver");

```

#### Procedure 18.6. Configure JdbcMixedStore Programmatically



1. Use the ***ConfigurationBuilder*** to create a new configuration object.
2. Add the ***JdbcMixedStore*** configuration builder to build a specific configuration related to this store.
3. The ***fetchPersistentState*** parameter determines whether or not to fetch the persistent state of a cache and apply it to the local cache store when joining the cluster. If the cache store is shared the fetch persistent state is ignored, as caches access the same cache store. A configuration exception will be thrown when starting the cache service if more than one cache loader has this property set to ***true***. The ***fetchPersistentState*** property is ***false*** by default.
4. The ***ignoreModifications*** parameter determines whether write methods are pushed to the specific cache loader by allowing write operations to the local file cache loader, but not the shared cache loader. In some cases, transient application data should only reside in a file-based cache loader on the same server as the in-memory cache. For example, this would apply with a further JDBC based cache loader used by all servers in the network. ***ignoreModifications*** is ***false*** by default.
5. The ***purgeOnStartup*** parameter specifies whether the cache is purged when initially started.
6. Configure the table as follows:
  - a. ***dropOnExit*** determines if the table will be dropped when the cache store is stopped. This is set to ***false*** by default.
  - b. ***createOnStart*** creates the table when starting the cache store if no table currently exists. This method is ***true*** by default.
  - c. ***tableNamePrefix*** sets the prefix for the name of the table in which the data will be stored.
  - d. The ***idColumnName*** property defines the column where the cache key or bucket ID is stored.
  - e. The ***dataColumnName*** property specifies the column where the cache entry or bucket is stored.
  - f. The ***timestampColumnName*** element specifies the column where the time stamp of the cache entry or bucket is stored.
7. The ***connectionPool*** element specifies a connection pool for the JDBC driver using the following parameters:
  - a. The ***connectionUrl*** parameter specifies the JDBC driver-specific connection URL.
  - b. The ***username*** parameter contains the username used to connect via the ***connectionUrl***.
  - c. The ***driverClass*** parameter specifies the class name of the driver used to connect to the database.



#### NOTE

Programmatic configurations can only be used with Red Hat JBoss Data Grid Library mode.

[Report a bug](#)

## 18.6.4. Cache Store Troubleshooting

### 18.6.4.1. IOExceptions with JdbcStringBasedStore

An `IOException` **Unsupported protocol version 48** error when using `JdbcStringBasedStore` indicates that your data column type is set to **VARCHAR**, **CLOB** or something similar instead of the correct type, **BLOB** or **VARBINARY**. Despite its name, `JdbcStringBasedStore` only requires that the keys are strings while the values can be any data type, so that they can be stored in a binary column.

[Report a bug](#)

## 18.7. THE REMOTE CACHE STORE

The `RemoteCacheStore` is an implementation of the cache loader that stores data in a remote Red Hat JBoss Data Grid cluster. The `RemoteCacheStore` uses the Hot Rod client-server architecture to communicate with the remote cluster.

For remote cache stores, Hot Rod provides load balancing, fault tolerance and the ability to fine tune the connection between the `RemoteCacheStore` and the cluster.

[Report a bug](#)

### 18.7.1. Remote Cache Store Configuration (Remote Client-Server Mode)

The following is a sample remote cache store configuration for Red Hat JBoss Data Grid's Remote Client-Server mode:

```
<remote-store cache="default"
              socket-timeout="60000"
              tcp-no-delay="true"
              hotrod-wrapping="true">
  <remote-server outbound-socket-binding="remote-store-hotrod-server" />
</remote-store>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, "Cache Store Configuration Details \(Remote Client-Server Mode\)"](#).

[Report a bug](#)

### 18.7.2. Remote Cache Store Configuration (Library Mode)

The following is a sample remote cache store configuration for Red Hat JBoss Data Grid's Library mode:

```
<persistence passivation="false">
  <remoteStore xmlns="urn:infinispan:config:remote:6.0"
              remoteCacheName="default"
              fetchPersistentState="false"
              shared="true"
              preload="false"
              ignoreModifications="false"
              purgeOnStartup="false">
```

```

        tcpNoDelay="true"
        pingOnStartup="true"
        keySizeEstimate="62"
        valueSizeEstimate="512"
        forceReturnValues="false">
<servers>
  <server host="127.0.0.1"
    port="19711"/>
</servers>
<connectionPool maxActive="99"
  maxIdle="97"
  maxTotal="98" />
</remoteStore>
</persistence>

```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).

[Report a bug](#)

### 18.7.3. Define the Outbound Socket for the Remote Cache Store

The Hot Rod server used by the remote cache store is defined using the **outbound-socket-binding** element in a **standalone.xml** file.

An example of this configuration in the **standalone.xml** file is as follows:

#### Example 18.1. Define the Outbound Socket

```

<server>
  <!-- Additional configuration elements here -->
  <socket-binding-group name="standard-sockets"
    default-interface="public"
    port-offset="{jboss.socket.binding.port-offset:0}">
    <!-- Additional configuration elements here -->
    <outbound-socket-binding name="remote-store-hotrod-server">
      <remote-destination host="remote-host"
        port="11222"/>
    </outbound-socket-binding>
  </socket-binding-group>
</server>

```

[Report a bug](#)

## 18.8. JPA CACHE STORE

The JPA (Java Persistence API) Cache Store stores cache entries in the database using a formal schema, which allows other applications to read the persisted data and load data provided by other applications into Red Hat JBoss Data Grid. The database should not be used by the other applications concurrently with JBoss Data Grid.



## IMPORTANT

In Red Hat JBoss Data Grid, JPA cache stores are only supported in Library mode.

[Report a bug](#)

### 18.8.1. JPA Cache Store Sample XML Configuration (Library Mode)

To configure JPA Cache Stores using XML in Red Hat JBoss Data Grid, add the following configuration to the `infinispan.xml` file:

```
<namedCache name="users">
  <!-- Insert additional configuration elements here -->
  <persistence passivation="false">
    <jpaStore xmlns="urn:infinispan:config:jpa:6.0"
      shared="true"
      preload="true"
      persistenceUnitName="MyPersistenceUnit"

entityClassName="org.infinispan.loaders.jpa.entity.User" />
  </persistence>
</namedCache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).

[Report a bug](#)

### 18.8.2. JPA Cache Store Sample Programmatic Configuration

To configure JPA Cache Stores programatically in Red Hat JBoss Data Grid, use the following:

```
Configuration cacheConfig = new
ConfigurationBuilder().persistence().addStore(JpaStoreConfigurationBuilder
.class).persistenceUnitName("org.infinispan.loaders.jpa.configurationTest"
)
    .entityClass(User.class)
    .build();
```

The parameters used in this code sample are as follows:

- The ***persistenceUnitName*** parameter specifies the name of the JPA cache store in the configuration file (`persistence.xml`) that contains the JPA entity class.
- The ***entityClass*** parameter specifies the JPA entity class that is stored in this cache. Only one class can be specified for each configuration.

[Report a bug](#)

### 18.8.3. Storing Metadata in the Database

When ***storeMetadata*** is set to **`true`** (default value), meta information about the entries such as expiration, creation and modification timestamps, and versioning is stored in the database. JBoss Data Grid stores the metadata in an additional table named `__ispn_metadata__` because the entity table

has a fixed layout that cannot accommodate the metadata.

The structure of this table depends on the database in use. Enable the automatic creation of this table using the same database as the test environment and then transfer the structure to the production database.

### Procedure 18.7. Configure `persistence.xml` for Metadata Entities

1. Using Hibernate as the JPA implementation allows automatic creation of these tables using the property `hibernate.hbm2ddl.auto` in `persistence.xml` as follows:

```
<property name="hibernate.hbm2ddl.auto" value="update"/>
```

2. Declare the metadata entity class to the JPA provider by adding the following to `persistence.xml`:

```
<class>org.infinispan.persistence.jpa.impl.MetadataEntity</class>
```

As outlined, metadata is always stored in a new table. If metadata information collection and storage is not required, set the `storeMetadata` attribute to `false` in the JPA Store configuration.

[Report a bug](#)

### 18.8.4. Deploying JPA Cache Stores in Various Containers

Red Hat JBoss Data Grid's JPA Cache Store implementations are deployed normally for all supported containers, except Red Hat JBoss Enterprise Application Platform. JBoss Data Grid's JBoss EAP modules contain the JPA cache store and related libraries (such as Hibernate). As a result, the relevant libraries are not packaged inside the application, but instead the application refers to the libraries in the JBoss EAP modules that have them installed.

These modules are not required for containers other than JBoss EAP. As a result, all the relevant libraries are packaged in the application's WAR/EAR file, such as with the following Maven dependency:

```
<dependency>
  <groupId>org.infinispan</groupId>
  <artifactId>infinispan-cachestore-jpa</artifactId>
  <version>6.3.0.Final-redhat-5</version>
</dependency>
```

### Procedure 18.8. Deploy JPA Cache Stores in JBoss EAP 6.3.x and earlier

- To add dependencies from the JBoss Data Grid modules to the application's classpath, provide the JBoss EAP deployer a list of dependencies in one of the following ways:
  - a. Add a dependency configuration to the `MANIFEST.MF` file:

```
Manifest-Version: 1.0
Dependencies: org.infinispan:jdg-6.5 services,
org.infinispan.persistence.jpa:jdg-6.5 services
```

- b. Add a dependency configuration to the `jboss-deployment-structure.xml` file:

```
<jboss-deployment-structure xmlns="urn:jboss:deployment-
structure:1.2">
  <deployment>
    <dependencies>
      <module name="org.infinispan.persistence.jpa"
slot="jdg-6.5" services="export"/>
      <module name="org.infinispan" slot="jdg-6.5"
services="export"/>
    </dependencies>
  </deployment>
</jboss-deployment-structure>
```

### Procedure 18.9. Deploy JPA Cache Stores in JBoss EAP 6.4 and later


1. Add the following property in **persistence.xml**:

```
<persistence-unit>
  [...]
  <properties>
    <property name="jboss.as.jpa.providerModule" value="application"
/>
  </properties>
</persistence-unit>
```

2. Add the following dependencies to the **jboss-deployment-structure.xml**:

```
<jboss-deployment-structure>
  <deployment>
    <dependencies>
      <module name="org.infinispan" slot="jdg-6.5"/>
      <module name="org.jgroups" slot="jdg-6.5"/>
      <module name="org.infinispan.persistence.jpa"
slot="jdg-6.5" services="export"/>
      <module name="org.hibernate"/>
    </dependencies>
  </deployment>
</jboss-deployment-structure>
```

3. Add any additional dependencies, such as additional JDG modules, are in use add these to the **dependencies** section in **jboss-deployment-structure.xml**.



#### IMPORTANT

JPA Cache Store is not supported in Apache Karaf in JBoss Data Grid 6.5.

[Report a bug](#)

## 18.9. CUSTOM CACHE STORES

Custom cache stores are a customized implementation of Red Hat JBoss Data Grid cache stores.

In order to create a custom cache store (or loader), implement all or a subset of the following interfaces based on the need:

- **CacheLoader**
- **CacheWriter**
- **AdvancedCacheLoader**
- **AdvancedCacheWriter**
- **ExternalStore**
- **AdvancedLoadWriteStore**

See [Section 17.1, “Cache Loaders and Cache Writers”](#) for individual functions of the interfaces.



#### NOTE

If the **AdvancedCacheWriter** is not implemented, the expired entries cannot be purged or cleared using the given writer.



#### NOTE

If the **AdvancedCacheLoader** is not implemented, the entries stored in the given loader will not be used for preloading and map/reduce iterations.

To migrate the existing cache store to the new API or to write a new store implementation, use **SingleFileStore** as an example. To view the **SingleFileStore** example code, download the JBoss Data Grid source code.

Use the following procedure to download **SingleFileStore** example code from the Customer Portal:

#### Procedure 18.10. Download JBoss Data Grid Source Code

1. To access the Red Hat Customer Portal, navigate to <https://access.redhat.com/home> in a browser.
2. Click **Downloads**.
3. In the box labeled **Red Hat JBoss Middleware**, click the **Download Software** button.
4. Enter the relevant credentials in the **Red Hat Login** and **Password** fields and click **Log In**.
5. In the **Software Downloads** page, select **Data Grid** from the list of drop down values.
6. From the list of downloadable files, locate **Red Hat JBoss Data Grid \${VERSION} Source Code** and click **Download**. Save and unpack it in a desired location.
7. Locate the **SingleFileStore** source code by navigating through `jboss-datagrid-6.5.1-sources/infinispan-6.3.1.Final-redhat-1-src/core/src/main/java/org/infinispan/persistence/file/SingleFileStore.java`.

[Report a bug](#)


### 18.9.1. Custom Cache Store Maven Archetype

An easy way to get started with developing a Custom Cache Store is to use the Maven archetype; creating an archetype will generate a new Maven project with the correct directory layout and sample code.

#### Procedure 18.11. Generate a Maven Archetype

1. Ensure the JBoss Data Grid Maven repository has been installed by following the instructions in the Red Hat JBoss Data Grid *Getting Started Guide*.
2. Open a command prompt and execute the following command to generate an archetype in the current directory:

```
mvn -Dmaven.repo.local="path/to/unzipped/jboss-datagrid-6.5.1-maven-repository/"
    archetype:generate
    -DarchetypeGroupId=org.infinispan
    -DarchetypeArtifactId=custom-cache-store-archetype
    -DarchetypeVersion=6.3.1.Final-redhat-1
```



#### NOTE

The above command has been broken into multiple lines for readability; however, when executed this command and all arguments must be on a single line.

[Report a bug](#)

### 18.9.2. Custom Cache Store Configuration (Remote Client-Server Mode)

The following is a sample configuration for a custom cache store in Red Hat JBoss Data Grid's Remote Client-Server mode:

#### Example 18.2. Custom Cache Store Configuration

```
<distributed-cache name="cacheStore" mode="SYNC" segments="20"
owners="2" remote-timeout="30000">
  <store class="my.package.CustomCacheStore">
    <properties>
      <property name="customStoreProperty" value="10" />
    </properties>
  </store>
</distributed-cache>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.3, "Cache Store Configuration Details \(Remote Client-Server Mode\)"](#).

[Report a bug](#)

#### 18.9.2.1. Option 1: Add Custom Cache Store using deployments (Remote Client-Server Mode)



**Procedure 18.12. Deploy Custom Cache Store .jar file to JDG server using deployments**

1. Add the following Java service loader file **META-INF/services/org.infinispan.persistence.spi.AdvancedLoadWriteStore** to the module and add a reference to the Custom Cache Store Class, such as seen below:

```
my.package.CustomCacheStore
```

2. Copy the jar to the **\$JDG\_HOME/standalone/deployments/** directory.
3. If the .jar file is available the server the following message will be displayed in the logs:

```
JBAS010287: Registering Deployed Cache Store service for store
'my.package.CustomCacheStore'
```

4. In the **infinispan-core** subsystem add an entry for the cache inside a **cache-container**, specifying the class that overrides one of the interfaces from [Section 18.9, “Custom Cache Stores”](#):

```
<subsystem xmlns="urn:infinispan:server:core:6.2">
  [...]
  <distributed-cache name="cacheStore" mode="SYNC" segments="20"
owners="2" remote-timeout="30000">
    <store class="my.package.CustomCacheStore">
      <!-- If custom properties are included these may be specified
as below -->
      <property name="customStoreProperty">10</property>
    </store>
  </distributed-cache>
  [...]
</subsystem>
```

[Report a bug](#)

**18.9.2.2. Option 2: Add Custom Cache Store using the CLI (Remote Client-Server Mode)****Procedure 18.13. Deploying Custom Cache Store .jar file to JDG server using the CLI**

1. Connect to the JDG server by running the below command:

```
$JDG_HOME] $ bin/jboss-cli.sh --connect=$IP:$PORT
```

2. Deploy the .jar file by executing the following command:

```
/] deploy /path/to/artifact.jar
```

[Report a bug](#)

**18.9.2.3. Option 3: Add Custom Cache Store using JON (Remote Client-Server Mode)****Procedure 18.14. Deploying Custom Cache Store .jar file to JDG server using JBoss Operation Network**

1. Log into JON.
2. Navigate to **Bundles** along the upper bar.
3. Click the **New** button and choose the **Recipe** radio button.
4. Insert a deployment bundle file content that references the store, similar to the following example:

```
<?xml version="1.0"?>
<project name="cc-bundle" default="main"
xmlns:rhq="antlib:org.rhq.bundle">

  <rhq:bundle name="Mongo DB Custom Cache Store" version="1.0"
description="Custom Cache Store">
    <rhq:deployment-unit name="JDG" compliance="full">
      <rhq:file name="custom-store.jar"/>
    </rhq:deployment-unit>
  </rhq:bundle>

  <target name="main" />
</project>
```

5. Proceed with **Next** button to **Bundle Groups** configuration wizard page and proceed with **Next** button once again.
6. Locate custom cache store **.jar** file using file uploader and **Upload** the file.
7. Proceed with **Next** button to **Summary** configuration wizard page. Proceed with **Finish** button in order to finish bundle configuration.
8. Navigate back to the **Bundles** tab along the upper bar.
9. Select the newly created bundle and click **Deploy** button.
10. Enter **Destination Name** and choose the proper Resource Group; this group should only consist of JDG servers.
11. Choose **Install Directory** from **Base Location's** radio box group.
12. Enter **/standalone/deployments** in **Deployment Directory** text field below.
13. Proceed with the wizard using the default options.
14. Validate the deployment using the following command on the server's host:

```
find $JDG_HOME -name "custom-store.jar"
```

15. Confirm the bundle has been installed in **\$JDG\_HOME/standalone/deployments**.

Once the above steps are completed the **.jar** file will be successfully uploaded and registered by the JDG server.

[Report a bug](#)

### 18.9.3. Custom Cache Store Configuration (Library Mode)

The following is a sample configuration for a custom cache store in Red Hat JBoss Data Grid's Library mode:

#### Example 18.3. Custom Cache Store Configuration

```
<persistence>
  <store class="org.infinispan.custom.CustomCacheStore"
    preload="true"
    shared="true">
    <properties>
      <property name="customStoreProperty"
        value="10" />
    </properties>
  </store>
</persistence>
```

For details about the elements and parameters used in this sample configuration, see [Section 18.2, “Cache Store Configuration Details \(Library Mode\)”](#).



#### NOTE

The Custom Cache Store classes must be in the classpath where Red Hat JBoss Data Grid is used. Most often this is accomplished by packaging the Custom Cache Store in with the application; however, it may also be accomplished by defining the Custom Cache Store as a module to EAP and listed as a dependency, as discussed in the Red Hat JBoss Enterprise Application Platform *Administration and Configuration Guide*.

[Report a bug](#)

## **PART VIII. SET UP PASSIVATION**

## CHAPTER 19. ACTIVATION AND PASSIVATION MODES

Activation is the process of loading an entry into memory and removing it from the cache store. Activation occurs when a thread attempts to access an entry that is in the store but not the memory (namely a passivated entry).

Passivation mode allows entries to be stored in the cache store after they are evicted from memory. Passivation prevents unnecessary and potentially expensive writes to the cache store. It is used for entries that are frequently used or referenced and therefore not evicted from memory.

While passivation is enabled, the cache store is used as an overflow tank, similar to virtual memory implementation in operating systems that swap memory pages to disk.

The passivation flag is used to toggle passivation mode, a mode that stores entries in the cache store only after they are evicted from memory.

[Report a bug](#)

### 19.1. PASSIVATION MODE BENEFITS

The primary benefit of passivation mode is that it prevents unnecessary and potentially expensive writes to the cache store. This is particularly useful if an entry is frequently used or referenced and therefore is not evicted from memory.

[Report a bug](#)

### 19.2. CONFIGURE PASSIVATION

In Red Hat JBoss Data Grid's Remote Client-Server mode, add the *passivation* parameter to the cache store element to toggle passivation for it:

#### Example 19.1. Toggle Passivation in Remote Client-Server Mode

```
<local-cache name="customCache"/>
  <!-- Additional configuration elements here -->
  <file-store passivation="true"
    <!-- Additional configuration elements here --> />
  <!-- Additional configuration elements here -->
</local-cache>
```

In Library mode, add the *passivation* parameter to the *persistence* element to toggle passivation:

#### Example 19.2. Toggle Passivation in Library Mode

```
<persistence passivation="true">
  <!-- Additional configuration elements here -->
</persistence>
```

[Report a bug](#)

## 19.3. EVICTION AND PASSIVATION

To ensure that a single copy of an entry remains, either in memory or in a cache store, use passivation in conjunction with eviction.

The primary reason to use passivation instead of a normal cache store is that updating entries require less resources when passivation is in use. This is because passivation does not require an update to the cache store.

[Report a bug](#)

### 19.3.1. Eviction and Passivation Usage

If the eviction policy caused the eviction of an entry from the cache while passivation is enabled, the following occur as a result:

- A notification regarding the passivated entry is emitted to the cache listeners.
- The evicted entry is stored.

When an attempt to retrieve an evicted entry is made, the entry is lazily loaded into memory from the cache loader. After the entry and its children are loaded, they are removed from the cache loader and a notification regarding the entry's activation is sent to the cache listeners.

[Report a bug](#)

### 19.3.2. Eviction Example when Passivation is Disabled

The following example indicates the state of the memory and the persistent store during eviction operations with passivation disabled.

**Table 19.1. Eviction when Passivation is Disabled**

Step	Key in Memory	Key on Disk
Insert <b>keyOne</b>	Memory: <b>keyOne</b>	Disk: <b>keyOne</b>
Insert <b>keyTwo</b>	Memory: <b>keyOne</b> , <b>keyTwo</b>	Disk: <b>keyOne</b> , <b>keyTwo</b>
Eviction thread runs, evicts <b>keyOne</b>	Memory: <b>keyTwo</b>	Disk: <b>keyOne</b> , <b>keyTwo</b>
Read <b>keyOne</b>	Memory: <b>keyOne</b> , <b>keyTwo</b>	Disk: <b>keyOne</b> , <b>keyTwo</b>
Eviction thread runs, evicts <b>keyTwo</b>	Memory: <b>keyOne</b>	Disk: <b>keyOne</b> , <b>keyTwo</b>
Remove <b>keyTwo</b>	Memory: <b>keyOne</b>	Disk: <b>keyOne</b>

[Report a bug](#)

### 19.3.3. Eviction Example when Passivation is Enabled

The following example indicates the state of the memory and the persistent store during eviction operations with passivation enabled.

**Table 19.2. Eviction when Passivation is Enabled**

Step	Key in Memory	Key on Disk
Insert <b>keyOne</b>	Memory: <b>keyOne</b>	Disk:
Insert <b>keyTwo</b>	Memory: <b>keyOne</b> , <b>keyTwo</b>	Disk:
Eviction thread runs, evicts <b>keyOne</b>	Memory: <b>keyTwo</b>	Disk: <b>keyOne</b>
Read <b>keyOne</b>	Memory: <b>keyOne</b> , <b>keyTwo</b>	Disk:
Eviction thread runs, evicts <b>keyTwo</b>	Memory: <b>keyOne</b>	Disk: <b>keyTwo</b>
Remove <b>keyTwo</b>	Memory: <b>keyOne</b>	Disk:

[Report a bug](#)

## **PART IX. SET UP CACHE WRITING**



## CHAPTER 20. CACHE WRITING MODES

Red Hat JBoss Data Grid presents configuration options with a single or multiple cache stores. This allows it to store data in a persistent location, for example a shared JDBC database or a local file system. JBoss Data Grid supports two caching modes:

- Write-Through (Synchronous)
- Write-Behind (Asynchronous)

[Report a bug](#)

### 20.1. WRITE-THROUGH CACHING

The Write-Through (or Synchronous) mode in Red Hat JBoss Data Grid ensures that when clients update a cache entry (usually via a `Cache.put()` invocation), the call does not return until JBoss Data Grid has located and updated the underlying cache store. This feature allows updates to the cache store to be concluded within the client thread boundaries.

[Report a bug](#)

#### 20.1.1. Write-Through Caching Benefits

The primary advantage of the Write-Through mode is that the cache and cache store are updated simultaneously, which ensures that the cache store remains consistent with the cache contents. This is at the cost of reduced performance for cache operations caused by the cache store accesses and updates during cache operations.

[Report a bug](#)

#### 20.1.2. Write-Through Caching Configuration (Library Mode)

No specific configuration operations are required to configure a Write-Through or synchronous cache store. All cache stores are Write-Through or synchronous unless explicitly marked as Write-Behind or asynchronous. The following procedure demonstrates a sample configuration file of a Write-Through unshared local file cache store.

##### Procedure 20.1. Configure a Write-Through Local File Cache Store

```
<?xml version="1.0" encoding="UTF-8"?>
<infinispan xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="urn:infinispan:config:6.2">
  <global />
  <default />
  <namedCache name="persistentCache">
    <persistence>
      <singleFile fetchPersistentState="true"
        ignoreModifications="false"
        purgeOnStartup="false"
        shared="false"
        location="${java.io.tmpdir}"/>
    </persistence>
  </namedCache>
</infinispan>
```

1. The ***name*** parameter specifies the name of the **`namedCache`** to use.
2. The ***fetchPersistentState*** parameter determines whether the persistent state is fetched when joining a cluster. Set this to **`true`** if using a replication and invalidation in a clustered environment. Additionally, if multiple cache stores are chained, only one cache store can have this property enabled. If a shared cache store is used, the cache does not allow a persistent state transfer despite this property being set to **`true`**. The ***fetchPersistentState*** parameter is **`false`** by default.
3. The ***ignoreModifications*** parameter determines whether operations that modify the cache (e.g. put, remove, clear, store, etc.) do not affect the cache store. As a result, the cache store can become out of sync with the cache.
4. The ***purgeOnStartup*** parameter specifies whether the cache is purged when initially started.
5. The ***shared*** parameter is used when multiple cache instances share a cache store and is now defined at the cache store level. This parameter can be set to prevent multiple cache instances writing the same modification multiple times. Valid values for this parameter are **`true`** and **`false`**.

[Report a bug](#)

## 20.2. WRITE-BEHIND CACHING

In Red Hat JBoss Data Grid's Write-Behind (Asynchronous) mode, cache updates are asynchronously written to the cache store. Asynchronous updates ensure that cache store updates are carried out by a thread different from the client thread interacting with the cache.

One of the foremost advantages of the Write-Behind mode is that the cache operation performance is not affected by the underlying store update. However, because of the asynchronous updates, for a brief period the cache store contains stale data compared to the cache.

[Report a bug](#)

### 20.2.1. About Unscheduled Write-Behind Strategy

In the Unscheduled Write-Behind Strategy mode, Red Hat JBoss Enterprise Data Grid attempts to store changes as quickly as possible by applying pending changes in parallel. This results in multiple threads waiting for modifications to conclude. Once these modifications are concluded, the threads become available and the modifications are applied to the underlying cache store.

This strategy is ideal for cache stores with low latency and low operational costs. An example of this is a local unshared file based cache store in which the cache store is local to the cache itself. Using this strategy the period of time where an inconsistency exists between the contents of the cache and the contents of the cache store is reduced to the shortest possible interval.

[Report a bug](#)

### 20.2.2. Unscheduled Write-Behind Strategy Configuration (Remote Client-Server Mode)

To set the write-behind strategy in Red Hat JBoss Data Grid's Remote Client-Server mode, add the **`write-behind`** element to the target cache store configuration as follows:

**Procedure 20.2. The write-behind Element**

```

<file-store passivation="false"
            path="{PATH}"
            purge="true"
            shared="false">
  <write-behind modification-queue-size="1024"
                shutdown-timeout="25000"
                flush-lock-timeout="15000"
                thread-pool-size="5" />
</file-store>

```

The **write-behind** element uses the following configuration parameters:

1. The **modification-queue-size** parameter sets the modification queue size for the asynchronous store. If updates occur faster than the cache store can process the queue, the asynchronous store behaves like a synchronous store. The store behavior remains synchronous and blocks elements until the queue is able to accept them, after which the store behavior becomes asynchronous again.
2. The **shutdown-timeout** parameter specifies the time in milliseconds after which the cache store is shut down. When the store is stopped some modifications may still need to be applied. Setting a large timeout value will reduce the chance of data loss. The default value for this parameter is **25000**.
3. The **flush-lock-timeout** parameter specifies the time (in milliseconds) to acquire the lock that guards the state to be periodically flushed. The default value for this parameter is **15000**.
4. The **thread-pool-size** parameter specifies the size of the thread pool. The threads in this thread pool apply modifications to the cache store. The default value for this parameter is **5**.

[Report a bug](#)

**20.2.3. Unscheduled Write-Behind Strategy Configuration (Library Mode)**

To enable the write-behind strategy of the cache entries to a store, add the **async** element to the store configuration as follows:

**Procedure 20.3. The async Element**

```

<persistence>
  <singleFile location="{LOCATION}">
    <async enabled="true"
           modificationQueueSize="1024"
           shutdownTimeout="25000"
           flushLockTimeout="15000"
           threadPoolSize="5"/>
  </singleFile>
</persistence>

```

The **async** element uses the following configuration parameters:

1. The **modificationQueueSize** parameter sets the modification queue size for the asynchronous store. If updates occur faster than the cache store can process the queue, the

asynchronous store behaves like a synchronous store. The store behavior remains synchronous and blocks elements until the queue is able to accept them, after which the store behavior becomes asynchronous again.

2. The ***shutdownTimeout*** parameter specifies the time in milliseconds after which the cache store is shut down. This provides time for the asynchronous writer to flush data to the store when a cache is shut down. The default value for this parameter is **25000**.
3. The ***flushLockTimeout*** parameter specifies the time (in milliseconds) to acquire the lock that guards the state to be periodically flushed. The default value for this parameter is **15000**.
4. The ***threadPoolSize*** parameter specifies the number of threads that concurrently apply modifications to the store. The default value for this parameter is **5**.

[Report a bug](#)

## **PART X. MONITOR CACHES AND CACHE MANAGERS**

## CHAPTER 21. SET UP JAVA MANAGEMENT EXTENSIONS (JMX)

### 21.1. ABOUT JAVA MANAGEMENT EXTENSIONS (JMX)

Java Management Extension (JMX) is a Java based technology that provides tools to manage and monitor applications, devices, system objects, and service oriented networks. Each of these objects is managed, and monitored by **MBeans**.

JMX is the de facto standard for middleware management and administration. As a result, JMX is used in Red Hat JBoss Data Grid to expose management and statistical information.

[Report a bug](#)

### 21.2. USING JMX WITH RED HAT JBOSS DATA GRID

Management in Red Hat JBoss Data Grid instances aims to expose as much relevant statistical information as possible. This information allows administrators to view the state of each instance. While a single installation can comprise of tens or hundreds of such instances, it is essential to expose and present the statistical information for each of them in a clear and concise manner.


In JBoss Data Grid, JMX is used in conjunction with JBoss Operations Network (JON) to expose this information and present it in an orderly and relevant manner to the administrator.

[Report a bug](#)

### 21.3. JMX STATISTIC LEVELS

JMX statistics can be enabled at two levels:

- At the cache level, where management information is generated by individual cache instances.
- At the **CacheManager** level, where the **CacheManager** is the entity that governs all cache instances created from it. As a result, the management information is generated for all these cache instances instead of individual caches.



#### IMPORTANT

In Red Hat JBoss Data Grid, statistics are enabled by default in Remote Client-Server mode and disabled by default for Library mode. While statistics are useful in assessing the status of JBoss Data Grid, they adversely affect performance and must be disabled if they are not required.

[Report a bug](#)

### 21.4. ENABLE JMX FOR CACHE INSTANCES

At the Cache level, JMX statistics can be enabled either declaratively or programmatically, as follows.

#### Enable JMX Declaratively at the Cache Level

Add the following snippet within either the `<default>` element for the default cache instance, or under the target `<namedCache>` element for a specific named cache:

```
<jmxStatistics enabled="true"/>
```

### Enable JMX Programmatically at the Cache Level

Add the following code to programmatically enable JMX at the cache level:

```
Configuration configuration = new
ConfigurationBuilder().jmxStatistics().enable().build();
```

[Report a bug](#)

## 21.5. ENABLE JMX FOR CACHEMANAGERS

At the **CacheManager** level, JMX statistics can be enabled either declaratively or programmatically, as follows.

### Enable JMX Declaratively at the CacheManager Level

Add the following in the `<global>` element to enable JMX declaratively at the **CacheManager** level:

```
<globalJmxStatistics enabled="true"/>
```

### Enable JMX Programmatically at the CacheManager Level

Add the following code to programmatically enable JMX at the **CacheManager** level:

```
GlobalConfiguration globalConfiguration = new
GlobalConfigurationBuilder().globalJmxStatistics().enable().build();
```

[Report a bug](#)

## 21.6. DISABLING THE CACHESTORE VIA JMX WHEN USING ROLLING UPGRADES

Red Hat JBoss Data Grid allows the `CacheStore` to be disabled via JMX by invoking the `disconnectSource` operation on the `RollingUpgradeManager` MBean.

### See Also:

- [Section C.17, “RollingUpgradeManager”](#)

[Report a bug](#)

## 21.7. MULTIPLE JMX DOMAINS

Multiple JMX domains are used when multiple **CacheManager** instances exist on a single virtual machine, or if the names of cache instances in different **CacheManagers** clash.

To resolve this issue, name each **CacheManager** in manner that allows it to be easily identified and used by monitoring tools such as JMX and JBoss Operations Network.

### Set a CacheManager Name Declaratively

Add the following snippet to the relevant **CacheManager** configuration:

```
<globalJmxStatistics enabled="true" cacheManagerName="Hibernate2LC"/>
```

### Set a CacheManager Name Programmatically

Add the following code to set the **CacheManager** name programmatically:

```
GlobalConfiguration globalConfiguration = new
GlobalConfigurationBuilder().globalJmxStatistics().enable().
cacheManagerName("Hibernate2LC").build();
```

[Report a bug](#)

## 21.8. MBEANS

An **MBean** represents a manageable resource such as a service, component, device or an application.

Red Hat JBoss Data Grid provides **MBeans** that monitor and manage multiple aspects. For example, **MBeans** that provide statistics on the transport layer are provided. If a JBoss Data Grid server is configured with JMX statistics, an **MBean** that provides information such as the hostname, port, bytes read, bytes written and the number of worker threads exists at the following location:

```
jboss.infinispan:type=Server, name=<Memcached|Hotrod>, component=Transport
```

**MBeans** are available under two JMX domains:

- `jboss.as` - these **MBeans** are created by the server subsystem.
- `jboss.infinispan` - these **MBeans** are symmetric to those created by embedded mode.

Only the **MBeans** under `jboss.infinispan` should be used for Red Hat JBoss Data Grid, as the ones under `jboss.as` are for Red Hat JBoss Enterprise Application Platform.



### NOTE

A full list of available MBeans, their supported operations and attributes, is available in the Appendix

[Report a bug](#)

### 21.8.1. Understanding MBeans

When JMX reporting is enabled at either the Cache Manager or Cache level, use a standard JMX GUI such as JConsole or VisualVM to connect to a Java Virtual Machine running Red Hat JBoss Data Grid. When connected, the following **MBeans** are available:

- If Cache Manager-level JMX statistics are enabled, an **MBean** named *`jboss.infinispan:type=CacheManager, name="DefaultCacheManager"`* exists, with properties specified by the Cache Manager **MBean**.
- If the cache-level JMX statistics are enabled, multiple **MBeans** display depending on the configuration in use. For example, if a write behind cache store is configured, an **MBean** that exposes properties that belong to the cache store component is displayed. All cache-level **MBeans** use the same format:



```
jboss.infinispan:type=Cache,name="<name-of-cache>( <cache-
mode>)",manager="<name-of-cache-manager>",component=<component-name>
```

In this format:

- Specify the default name for the cache using the **cache-container** element's **default-cache** attribute.
- The **cache-mode** is replaced by the cache mode of the cache. The lower case version of the possible enumeration values represents the cache mode.
- The **component-name** is replaced by one of the JMX component names from the JMX reference documentation.

As an example, the cache store JMX component **MBean** for a default cache configured for synchronous distribution would be named as follows:

```
jboss.infinispan:type=Cache,name="default(dist_sync)",
manager="default",component=CacheStore
```

Each cache and cache manager name is within quotation marks to prevent the use of unsupported characters in these user-defined names.

[Report a bug](#)

## 21.8.2. Registering MBeans in Non-Default MBean Servers

The default location where all the MBeans used are registered is the standard JVM MBeanServer platform. Users can set up an alternative MBeanServer instance as well. Implement the MBeanServerLookup interface to ensure that the **getMBeanServer()** method returns the desired (non default) MBeanServer.

To set up a non default location to register your MBeans, create the implementation and then configure Red Hat JBoss Data Grid with the fully qualified name of the class. An example is as follows:

### To Add the Fully Qualified Domain Name Declaratively

Add the following snippet:

```
<globalJmxStatistics enabled="true"
mBeanServerLookup="com.acme.MyMBeanServerLookup"/>
```

### To Add the Fully Qualified Domain Name Programmatically

Add the following code:


```
GlobalConfiguration globalConfiguration = new
GlobalConfigurationBuilder().globalJmxStatistics().enable().
mBeanServerLookup("com.acme.MyMBeanServerLookup").build();
```

[Report a bug](#)

## CHAPTER 22. SET UP JBOSS OPERATIONS NETWORK (JON)


### 22.1. ABOUT JBOSS OPERATIONS NETWORK (JON)

The JBoss Operations Network (JON) is JBoss' administration and management platform used to develop, test, deploy and monitor the application life cycle. JBoss Operations Network is JBoss' enterprise management solution and is recommended for the management of multiple Red Hat JBoss Data Grid instances across servers. JBoss Operations Network's agent and auto discovery features facilitate monitoring the Cache Manager and Cache instances in JBoss Data Grid. JBoss Operations Network presents graphical views of key runtime parameters and statistics and allows administrators to set thresholds and be notified if usage exceeds or falls under the set thresholds.



#### IMPORTANT

In Red Hat JBoss Data Grid Remote Client-Server mode, statistics are enabled by default. While statistics are useful in assessing the status of JBoss Data Grid, they adversely affect performance and must be disabled if they are not required. In JBoss Data Grid Library mode, statistics are disabled by default and must be explicitly enabled when required.



#### IMPORTANT

To achieve full functionality of JBoss Operations Network library plugin for JBoss Data Grid's Library mode, upgrade to JBoss Operations Network 3.2.0 with patch **Update 02** or higher version. For information on upgrading the JBoss Operations Network, see the *Upgrading JBoss ON* section in the JBoss Operations Network *Installation Guide*.

[Report a bug](#)

### 22.2. DOWNLOAD JBOSS OPERATIONS NETWORK (JON)

#### 22.2.1. Prerequisites for Installing JBoss Operations Network (JON)

In order to install JBoss Operations Network in Red Hat JBoss Data Grid, the following is required:

- A Linux, Windows, or Mac OSX operating system, and an x86\_64, i686, or ia64 processor.
- Java 6 or higher is required to run both the JBoss Operations Network Server and the JBoss Operations Network Agent.
- Synchronized clocks on JBoss Operations Network Servers and Agents.
- An external database must be installed.

[Report a bug](#)

#### 22.2.2. Download JBoss Operations Network

Use the following procedure to download Red Hat JBoss Operations Network (JON) from the Customer Portal:

### Procedure 22.1. Download JBoss Operations Network

1. To access the Red Hat Customer Portal, navigate to <https://access.redhat.com/home> in a browser.
2. Click **Downloads**.
3. In the box labeled **Red Hat JBoss Middleware**, click the **Download Software** button.
4. Enter the relevant credentials in the **Red Hat Login** and **Password** fields and click **Log In**.
5. In the **Software Downloads** page, select **JBoss Operations Network** from the list of drop down values.
6. Select the appropriate version in the **Version** drop down menu list.
7. Click the **Download** button next to the desired download file.

[Report a bug](#)

### 22.2.3. Remote JMX Port Values

A port value must be provided to allow Red Hat JBoss Data Grid instances to be located. The value itself can be any available port.

Provide unique (and available) remote JMX ports to run multiple JBoss Data Grid instances on a single machine. A locally running JBoss Operations Network agent can discover each instance using the remote port values.

[Report a bug](#)

### 22.2.4. Download JBoss Operations Network (JON) Plugin

Complete this task to download the JBoss Operations Network (JON) plugin for Red Hat JBoss Data Grid from the Red Hat Customer Portal.

### Procedure 22.2. Download Installation Files

1. Open <http://access.redhat.com> in a web browser.
2. Click **Downloads** in the menu across the top of the page.
3. Click **Downloads** in the list under JBoss Enterprise Middleware.
4. Enter your login information.

You are taken to the Software Downloads page.

5. **Download the JBoss Operations Network Plugin**

If you intend to use the JBoss Operations Network plugin for JBoss Data Grid, select **JBoss ON for JDG** from either the Software Downloads drop-down box, or the menu on the left.


- a. Click the **JBoss Operations Network VERSION Base Distribution** download link.
- b. Click the **Download** link to start the Base Distribution download.

- c. Repeat the steps to download the **JDG Plugin Pack for JBoss ON VERSION**

[Report a bug](#)

## 22.3. JBOSS OPERATIONS NETWORK SERVER INSTALLATION

The core of JBoss Operations Network is the server, which communicates with agents, maintains the inventory, manages resource settings, interacts with content providers, and provides a central management UI.



### NOTE

For more detailed information about configuring JBoss Operations Network, see the JBoss Operations Network *Installation Guide*.

[Report a bug](#)

## 22.4. JBOSS OPERATIONS NETWORK AGENT

The JBoss Operations Network Agent is a standalone Java application. Only one agent is required per machine, regardless of how many resources you require the agent to manage.

The JBoss Operations Network Agent does not ship fully configured. Once the agent has been installed and configured it can be run as a Windows service from a console, or run as a daemon or **init.d** script in a UNIX environment.

A JBoss Operations Network Agent must be installed on each of the machines being monitored in order to collect data.

The JBoss Operations Network Agent is typically installed on the same machine on which Red Hat JBoss Data Grid is running, however where there are multiple machines an agent must be installed on each machine.



### NOTE

For more detailed information about configuring JBoss Operations Network agents, see the JBoss Operations Network *Installation Guide*.

[Report a bug](#)

## 22.5. JBOSS OPERATIONS NETWORK FOR REMOTE CLIENT-SERVER MODE

In Red Hat JBoss Data Grid's Remote Client-Server mode, the JBoss Operations Network plug-in is used to

- initiate and perform installation and configuration operations.
- monitor resources and their metrics.

In Remote Client-Server mode, the JBoss Operations Network plug-in uses JBoss Enterprise Application Platform's management protocol to obtain metrics and perform operations on the JBoss Data Grid server.

[Report a bug](#)

## 22.5.1. Installing the JBoss Operations Network Plug-in (Remote Client-Server Mode)

The following procedure details how to install the JBoss Operations Network plug-ins for Red Hat JBoss Data Grid's Remote Client-Server mode.

### 1. Install the plug-ins

- Copy the JBoss Data Grid server rhq plug-in to `$JON_SERVER_HOME/plugins`.
- Copy the JBoss Enterprise Application Platform plug-in to `$JON_SERVER_HOME/plugins`.

The server will automatically discover plug-ins here and deploy them. The plug-ins will be removed from the plug-ins directory after successful deployment.

### 2. Obtain plug-ins

Obtain all available plug-ins from the JBoss Operations Network server. To do this, type the following into the agent's console:

```
plugins update
```

### 3. List installed plug-ins

Ensure the JBoss Enterprise Application Platform plug-in and the JBoss Data Grid server rhq plug-in are installed correctly using the following:

```
plugins info
```

JBoss Operation Network can now discover running JBoss Data Grid servers.

[Report a bug](#)

## 22.6. JBOSS OPERATIONS NETWORK REMOTE-CLIENT SERVER PLUGIN

### 22.6.1. JBoss Operations Network Plugin Metrics

Table 22.1. JBoss Operations Network Traits for the Cache Container (Cache Manager)

Trait Name	Display Name	Description
cache-manager-status	Cache Container Status	The current runtime status of a cache container.
cluster-name	Cluster Name	The name of the cluster.
members	Cluster Members	The names of the members of the cluster.
coordinator-address	Coordinator Address	The coordinator node's address.

Trait Name	Display Name	Description
local-address	Local Address	The local node's address.
version	Version	The cache manager version.
defined-cache-names	Defined Cache Names	The caches that have been defined for this manager.

**Table 22.2. JBoss Operations Network Metrics for the Cache Container (Cache Manager)**

Metric Name	Display Name	Description
cluster-size	Cluster Size	How many members are in the cluster.
defined-cache-count	Defined Cache Count	How many caches that have been defined for this manager.
running-cache-count	Running Cache Count	How many caches are running under this manager.
created-cache-count	Created Cache Count	How many caches have actually been created under this manager.

**Table 22.3. JBoss Operations Network Traits for the Cache**


Trait Name	Display Name	Description
cache-status	Cache Status	The current runtime status of a cache.
cache-name	Cache Name	The current name of the cache.
version	Version	The cache version.

**Table 22.4. JBoss Operations Network Metrics for the Cache**

Metric Name	Display Name	Description
cache-status	Cache Status	The current runtime status of a cache.
number-of-locks-available	[LockManager] Number of locks available	The number of exclusive locks that are currently available.

Metric Name	Display Name	Description
concurrency-level	[LockManager] Concurrency level	The LockManager's configured concurrency level.
average-read-time	[Statistics] Average read time	Average number of milliseconds required for a read operation on the cache to complete.
hit-ratio	[Statistics] Hit ratio	The result (in percentage) when the number of hits (successful attempts) is divided by the total number of attempts.
elapsed-time	[Statistics] Seconds since cache started	The number of seconds since the cache started.
read-write-ratio	[Statistics] Read/write ratio	The read/write ratio (in percentage) for the cache.
average-write-time	[Statistics] Average write time	Average number of milliseconds a write operation on a cache requires to complete.
hits	[Statistics] Number of cache hits	Number of cache hits.
evictions	[Statistics] Number of cache evictions	Number of cache eviction operations.
remove-misses	[Statistics] Number of cache removal misses	Number of cache removals where the key was not found.
time-since-reset	[Statistics] Seconds since cache statistics were reset	Number of seconds since the last cache statistics reset.
number-of-entries	[Statistics] Number of current cache entries	Number of entries currently in the cache.
stores	[Statistics] Number of cache puts	Number of cache put operations
remove-hits	[Statistics] Number of cache removal hits	Number of cache removal operation hits.
misses	[Statistics] Number of cache misses	Number of cache misses.
success-ratio	[RpcManager] Successful replication ratio	Successful replications as a ratio of total replications in numeric double format.

Metric Name	Display Name	Description
replication-count	[RpcManager] Number of successful replications	Number of successful replications
replication-failures	[RpcManager] Number of failed replications	Number of failed replications
average-replication-time	[RpcManager] Average time spent in the transport layer	The average time (in milliseconds) spent in the transport layer.
commits	[Transactions] Commits	Number of transaction commits performed since the last reset.
prepares	[Transactions] Prepares	Number of transaction prepares performed since the last reset.
rollbacks	[Transactions] Rollbacks	Number of transaction rollbacks performed since the last reset.
invalidations	[Invalidation] Number of invalidations	Number of invalidations.
passivations	[Passivation] Number of cache passivations	Number of passivation events.
activations	[Activations] Number of cache entries activated	Number of activation events.
cache-loader-loads	[Activation] Number of cache store loads	Number of entries loaded from the cache store.
cache-loader-misses	[Activation] Number of cache store misses	Number of entries that did not exist in the cache store.
cache-loader-stores	[CacheStore] Number of cache store stores	Number of entries stored in the cache stores.

**NOTE**

Gathering of some of these statistics is disabled by default.

**JBoss Operations Network Metrics for Connectors**

The metrics provided by the JBoss Operations Network (JON) plugin for Red Hat JBoss Data Grid are for REST and Hot Rod endpoints only. For the REST protocol, the data must be taken from the Web subsystem metrics. For details about each of these endpoints, see the *Getting Started Guide*.

**Table 22.5. JBoss Operations Network Metrics for the Connectors**



Metric Name	Display Name	Description
bytesRead	Bytes Read	Number of bytes read.
bytesWritten	Bytes Written	Number of bytes written.

**NOTE**

Gathering of these statistics is disabled by default.

[Report a bug](#)

## 22.6.2. JBoss Operations Network Plugin Operations

**Table 22.6. JBoss ON Plugin Operations for the Cache**

Operation Name	Description
Start Cache	Starts the cache.
Stop Cache	Stops the cache.
Clear Cache	Clears the cache contents.
Reset Statistics	Resets statistics gathered by the cache.
Reset Activation Statistics	Resets activation statistics gathered by the cache.
Reset Invalidation Statistics	Resets invalidations statistics gathered by the cache.
Reset Passivation Statistics	Resets passivation statistics gathered by the cache.
Reset Rpc Statistics	Resets replication statistics gathered by the cache.
Remove Cache	Removes the given cache from the cache-container.
Record Known Global Keyset	Records the global known keyset to a well-known key for retrieval by the upgrade process.
Synchronize Data	Synchronizes data from the old cluster to this using the specified migrator.
Disconnect Source	Disconnects the target cluster from the source cluster according to the specified migrator.

### JBoss Operations Network Plugin Operations for the Cache Backups

The cache backups used for these operations are configured using cross-datacenter replication. In the JBoss Operations Network (JON) User Interface, each cache backup is the child of a cache. For more information about cross-datacenter replication, see [Chapter 30, Set Up Cross-Datacenter Replication](#)

**Table 22.7. JBoss Operations Network Plugin Operations for the Cache Backups**

Operation Name	Description
status	Display the site status.
bring-site-online	Brings the site online.
take-site-offline	Takes the site offline.

### Cache (Transactions)

Red Hat JBoss Data Grid does not support using Transactions in Remote Client-Server mode. As a result, none of the endpoints can use transactions.

[Report a bug](#)

## 22.6.3. JBoss Operations Network Plugin Attributes

**Table 22.8. JBoss ON Plugin Attributes for the Cache (Transport)**

Attribute Name	Type	Description
cluster	string	The name of the group communication cluster.
executor	string	The executor used for the transport.
lock-timeout	long	The timeout period for locks on the transport. The default value is <b>240000</b> .
machine	string	A machine identifier for the transport.
rack	string	A rack identifier for the transport.
site	string	A site identifier for the transport.
stack	string	The JGroups stack used for the transport.

[Report a bug](#)

## 22.6.4. Create a New Cache Using JBoss Operations Network (JON)

Use the following steps to create a new cache using JBoss Operations Network (JON) for Remote Client-Server mode.

### Procedure 22.3. Creating a new cache in Remote Client-Server mode

1. Log into the JBoss Operations Network Console.
  - a. From from the JBoss Operations Network console, click **Inventory**.
  - b. Select **Servers** from the **Resources** list on the left of the console.
2. Select the specific Red Hat JBoss Data Grid server from the servers list.
  - a. Below the server name, click **infinispan** and then **Cache Containers**.
3. Select the desired cache container that will be parent for the newly created cache.
  - a. Right-click the selected cache container. For example, **clustered**.
  - b. In the context menu, navigate to **Create Child** and select **Cache**.
4. Create a new cache in the resource create wizard.
  - a. Enter the new cache name and click **Next**.
  - b. Set the cache attributes in the Deployment Options and click **Finish**.



#### NOTE

Refresh the view of caches in order to see newly added resource. It may take several minutes for the Resource to show up in the Inventory.

[Report a bug](#)

## 22.7. JBOSS OPERATIONS NETWORK FOR LIBRARY MODE

In Red Hat JBoss Data Grid's Library mode, the JBoss Operations Network plug-in is used to

- initiate and perform installation and configuration operations.
- monitor resources and their metrics.

In Library mode, the JBoss Operations Network plug-in uses JMX to obtain metrics and perform operations on an application using the JBoss Data Grid library.

[Report a bug](#)

### 22.7.1. Installing the JBoss Operations Network Plug-in (Library Mode)

Use the following procedure to install the JBoss Operations Network plug-in for Red Hat JBoss Data Grid's Library mode.

#### Procedure 22.4. Install JBoss Operations Network Library Mode Plug-in

1. **Open the JBoss Operations Network Console**

- a. From the JBoss Operations Network console, select **Administration**.
- b. Select **Agent Plugins** from the **Configuration** options on the left side of the console.

Name	Description	Last Updated	Enabled?	Deployed?
<b>Abstract Augeas Plugin</b>	An abstract plugin supporting concrete plugins that are based on Augeas, the open-source configuration library for Linux	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Abstract Database</b>	Abstract plugin supporting concrete database plugins	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Ant Bundle Plugin</b>		Sep 20, 2012 1:35:12 PM	✓	✓
<b>Apache HTTP Server</b>	Management of Apache web servers	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Generic JMX</b>	Supports management of JMX MBean Servers via various remoting systems.	Sep 20, 2012 1:35:12 PM	✓	✓
<b>IIS</b>	Monitoring of Microsoft IIS Services	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Operating System Platforms</b>	Management and monitoring of operating system level functions	Sep 20, 2012 1:35:12 PM	✓	✓
<b>PostgreSQL Database</b>	Management and monitoring of PostgreSQL versions 8.2 and higher	Sep 20, 2012 1:35:12 PM	✓	✓
<b>RHQ Agent</b>	Management and monitoring of the RHQ Agent	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Script</b>	Manages resources that have command line executables or scripts as their management interfaces	Sep 20, 2012 1:35:12 PM	✓	✓

Figure 22.1. JBoss Operations Network Console for JBoss Data Grid

## 2. Upload the Library Mode Plug-in

- a. Click **Browse**, locate the **InfinispanPlugin** on your local file system.
- b. Click **Upload** to add the plug-in to the JBoss Operations Network Server.

File successfully uploaded

Name ^	Description	Last Updated	Enabled?	Deployed?
<b>Abstract Augeas Plugin</b>	An abstract plugin supporting concrete plugins that are based on Augeas, the open-source configuration library for Linux	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Abstract Database</b>	Abstract plugin supporting concrete database plugins	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Ant Bundle Plugin</b>		Sep 20, 2012 1:35:12 PM	✓	✓
<b>Apache HTTP Server</b>	Management of Apache web servers	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Generic JMX</b>	Supports management of JMX MBean Servers via various remoting systems.	Sep 20, 2012 1:35:12 PM	✓	✓
<b>IIS</b>	Monitoring of Microsoft IIS Services	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Operating System Platforms</b>	Management and monitoring of operating system level functions	Sep 20, 2012 1:35:12 PM	✓	✓
<b>PostgreSQL Database</b>	Management and monitoring of PostgreSQL versions 8.2 and higher	Sep 20, 2012 1:35:12 PM	✓	✓
<b>RHQ Agent</b>	Management and monitoring of the RHQ Agent	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Script</b>	Manages resources that have command line executables or scripts as their management interfaces	Sep 20, 2012 1:35:12 PM	✓	✓

Scan For Updates Hide Deleted Upload Plugin :  Browse... Upload ✓

Enable Disable Delete Purge Refresh Total Rows: 10 (selected: 0)

Figure 22.2. Upload the InfinispanPlugin.

### 3. Scan for Updates

- a. Once the file has successfully uploaded, click **Scan For Updates** at the bottom of the screen.
- b. The **InfinispanPlugin** will now appear in the list of installed plug-ins.

Name ^	Description	Last Updated	Enabled?	Deployed?
<b>Abstract Augeas Plugin</b>	An abstract plugin supporting concrete plugins that are based on Augeas, the open-source configuration library for Linux	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Abstract Database</b>	Abstract plugin supporting concrete database plugins	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Ant Bundle Plugin</b>		Sep 20, 2012 1:35:12 PM	✓	✓
<b>Apache HTTP Server</b>	Management of Apache web servers	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Generic JMX</b>	Supports management of JMX MBean Servers via various remoting systems.	Sep 20, 2012 1:35:12 PM	✓	✓
<b>IIS</b>	Monitoring of Microsoft IIS Services	Sep 20, 2012 1:35:12 PM	✓	✓
<b>InfinispanPlugin</b>	Supports management and monitoring of Infinispan	Nov 12, 2012 11:39:25 AM	✓	✓
<b>Operating System Platforms</b>	Management and monitoring of operating system level functions	Sep 20, 2012 1:35:12 PM	✓	✓
<b>PostgreSQL Database</b>	Management and monitoring of PostgreSQL versions 8.2 and higher	Sep 20, 2012 1:35:12 PM	✓	✓
<b>RHQ Agent</b>	Management and monitoring of the RHO Agent	Sep 20, 2012 1:35:12 PM	✓	✓
<b>Script</b>	Manages resources that have command line executables or scripts as their management interfaces	Sep 20, 2012 1:35:12 PM	✓	✓

Figure 22.3. Scan for Updated Plug-ins.

[Report a bug](#)

## 22.7.2. Monitoring Of JBoss Data Grid Instances in Library Mode

### 22.7.2.1. Prerequisites

The following is a list of common prerequisites for [Section 22.7.2.3.1, “Monitor an Application Deployed in Standalone Mode”](#), [Section 22.7.2.3.2, “Monitor an Application Deployed in Domain Mode”](#) and [Section 22.7.2.2.1, “Manually Adding JBoss Data Grid Instances in Library Mode”](#).

- A correctly configured instance of JBoss Operations Network (JON) 3.2.0 with patch **Update 02** or higher version.
- A running instance of JON Agent on the server where the application will run. For more information, see [Section 22.4, “JBoss Operations Network Agent”](#)
- An operational instance of the RHO agent with a full JDK. Ensure that the agent has access to the **tools.jar** file from the JDK in particular. In the JON agent's environment file (**bin/rhq-env.sh**), set the value of the **RHQ\_AGENT\_JAVA\_HOME** property to point to a full JDK home.
- The RHO agent must have been initiated using the same user as the JBoss Enterprise Application Platform instance. As an example, running the JON agent as a user with root privileges and the JBoss Enterprise Application Platform process under a different user does not work as expected and must be avoided.
- An installed JON plugin for JBoss Data Grid Library Mode. For more information, see [Section 22.7.1, “Installing the JBoss Operations Network Plug-in \(Library Mode\)”](#)

- **Generic JMX plugin** from JBoss Operation Networks 3.2.0 with patch **Update 02** or better version in use.
- A custom application using Red Hat JBoss Data Grid's Library mode with enabled JMX statistics for library mode caches in order to make statistics and monitoring working. For details how to enable JMX statistics for cache instances, see [Section 21.4, “Enable JMX for Cache Instances”](#) and to enable JMX for cache managers see [Section 21.5, “Enable JMX for CacheManagers”](#)
- The Java Virtual Machine (JVM) must be configured to expose the JMX MBean Server. For the Oracle/Sun JDK, see <http://docs.oracle.com/javase/1.5.0/docs/guide/management/agent.html>
- A correctly added and configured management user for JBoss Enterprise Application Platform.

[Report a bug](#)

## 22.7.2.2. Manually Adding JBoss Data Grid Instances in Library Mode

### 22.7.2.2.1. Manually Adding JBoss Data Grid Instances in Library Mode

To add Red Hat JBoss Data Grid instances to JBoss Operations Network manually, use the following procedure in the JBoss Operations Network interface.

#### Procedure 22.5. Add JBoss Data Grid Instances in Library Mode

##### 1. Import the Platform

- a. Navigate to the **Inventory** and select **Discovery Queue** from the **Resources** list on the left of the console.
- b. Select the platform on which the application is running and click **Import** at the bottom of the screen.

The screenshot shows the JBoss Administration Console interface. The top navigation bar includes 'Dashboard', 'Inventory', 'Reports', 'Bundles', 'Administration', and 'Help'. The user is logged in as 'rhqadmin'. The left sidebar shows a tree view under 'Resources' with 'Discovery Queue' selected, and 'Groups' with 'All Groups' selected. The main content area is titled 'Discovery Queue' and contains a table with the following data:

Resource Name	Resource Key	Resource Type	Description	Inventory Status	Discovery Time
jdg	jdg	Linux	Linux Operating System	New	Nov 12, 2012 11:40:55 AM

At the bottom of the table, there are buttons for 'Import', 'Ignore', 'Unignore', a 'Show' dropdown menu set to 'New', 'Select All', and 'Deselect All'.

Figure 22.4. Import the Platform from the Discovery Queue.

## 2. Access the Servers on the Platform

- a. The **jdg** Platform now appears in the **Plat forms** list.
- b. Click on the Platform to access the servers that are running on it.



The screenshot shows the JBoss Operations Network (JON) interface. The top navigation bar includes 'Dashboard', 'Inventory' (selected), 'Reports', 'Bundles', 'Administration', and 'Help'. The user is logged in as 'rhqadmin'. The left sidebar shows a tree view under 'Resources' with 'Platforms' selected. The main content area displays the 'Platforms' section with a search bar and a table. The table has columns for Name, Ancestry, Description, Type, Version, and Availability. One row is visible for the 'jdg' platform, which is a 'Linux Operating System' of type 'Linux' with version 'Linux 3.5.0-18-generic' and is available (indicated by a green checkmark). At the bottom of the table, there are buttons for 'Uninventory', 'Disable', 'Enable', and 'Refresh', along with a status indicator 'Total Rows: 1 (selected: 0)'.

Name	Ancestry	Description	Type	Version	Availability
jdg		Linux Operating System	Linux	Linux 3.5.0-18-generic	✓

Figure 22.5. Open the `jdg` Platform to view the list of servers.

### 3. Import the JMX Server

- a. From the **Inventory** tab, select **Child Resources**.
- b. Click the **Import** button at the bottom of the screen and select the **JMX Server** option from the list.






Figure 22.7. Select the JDK 5 Template.

## 5. Modify the Connector Address

- a. In the **Deployment Options** menu, modify the supplied **Connector Address** with the hostname and JMX port of the process containing the Infinispan Library.
- b. Enter the JMX connector address of the new JBoss Data Grid instance you want to monitor. For example:

Connector Address:

```
service:jmx:rmi:///127.0.0.1/jndi/rmi:///127.0.0.1:7997/jmxrmi
```



### NOTE

The connector address varies depending on the host and the JMX port assigned to the new instance. In this case, instances require the following system properties at start up:

```
-Dcom.sun.management.jmxremote.port=7997 -
Dcom.sun.management.jmxremote.ssl=false -
Dcom.sun.management.jmxremote.authenticate=false
```

- c. Specify the **Principal** and **Credentials** information if required.
- d. Click **Finish**.

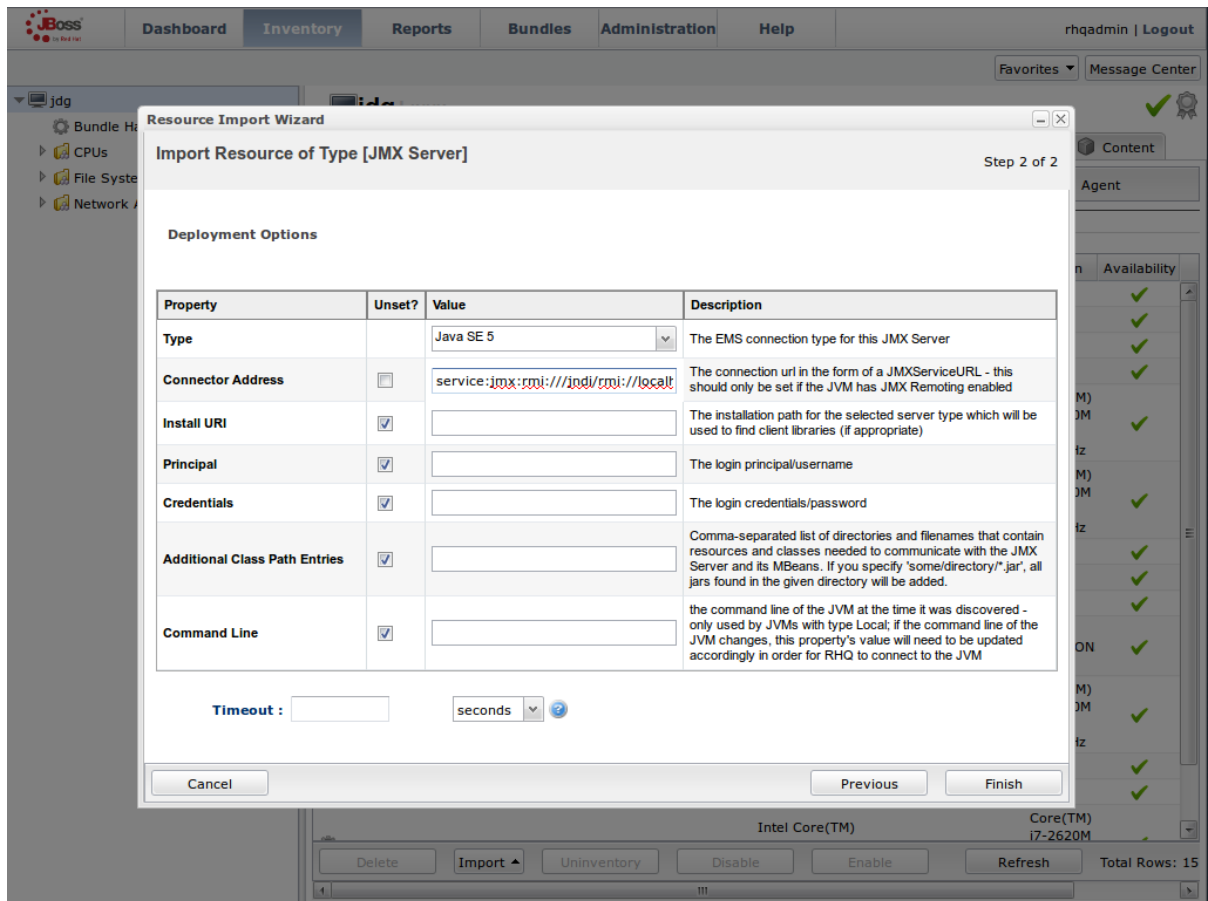


Figure 22.8. Modify the values in the Deployment Options screen.

## 6. View Cache Statistics and Operations

- a. Click **Refresh** to refresh the list of servers.
- b. The **JMX Servers** tree in the panel on the left side of the screen contains the **Infinispan Cache Managers** node, which contains the available cache managers. The available cache managers contain the available caches.
- c. Select a cache from the available caches to view metrics.
- d. Select the **Monitoring** tab.
- e. The **Tables** view shows statistics and metrics.
- f. The **Operations** tab provides access to the various operations that can be performed on the services.

Name	Alerts	Minimum	Maximum	Average	Last
[RpcManager] Average time spent in the transport layer	0	NaN	NaN	NaN	NaN
[RpcManager] Number of failed replications	0	NaN	NaN	NaN	NaN
[RpcManager] Number of successful replications	0	NaN	NaN	NaN	NaN
[RpcManager] Successful replication ratio	0	NaN	NaN	NaN	NaN
[Statistics] Average read time	0	NaN	NaN	NaN	0
[Statistics] Average write time	0	NaN	NaN	NaN	0
[Statistics] Hit ratio	0	NaN	NaN	NaN	0.8333333333333333
[Statistics] Number of cache evictions	0	NaN	NaN	NaN	0
[Statistics] Number of cache hits	0	NaN	NaN	NaN	5
[Statistics] Number of cache misses	0	NaN	NaN	NaN	1
[Statistics] Number of cache puts	0	NaN	NaN	NaN	2
[Statistics] Number of cache removal hits	0	NaN	NaN	NaN	0
[Statistics] Number of cache removal misses	0	NaN	NaN	NaN	0
[Statistics] Number of current cache entries	0	NaN	NaN	NaN	2
[Statistics] Read/write ratio	0	NaN	NaN	NaN	3
[Statistics] Seconds since cache started	0	NaN	NaN	NaN	7376
[Statistics] Seconds since cache	0	NaN	NaN	NaN	7376

Figure 22.9. Metrics and operational data relayed through JMX is now available in the JBoss Operations Network console.

[Report a bug](#)

### 22.7.2.3. Monitor Custom Applications Using Library Mode Deployed On JBoss Enterprise Application Platform

#### 22.7.2.3.1. Monitor an Application Deployed in Standalone Mode

Use the following instructions to monitor an application deployed in JBoss Enterprise Application Platform using its standalone mode:

#### Procedure 22.6. Monitor an Application Deployed in Standalone Mode

##### 1. Start the JBoss Enterprise Application Platform Instance

Start the JBoss Enterprise Application Platform instance as follows:

- a. Enter the following command at the command line or change standalone configuration file (`/bin/standalone.conf`) respectively:

```
JAVA_OPTS="$JAVA_OPTS -Dorg.rhq.resourceKey=MyEAP"
```

- b. Start the JBoss Enterprise Application Platform instance in standalone mode as follows:

```
$JBOSS_HOME/bin/standalone.sh
```

##### 2. Deploy the Red Hat JBoss Data Grid Application

Deploy the WAR file that contains the JBoss Data Grid Library mode application with `globalJmxStatistics` and `jmxStatistics` enabled.

### 3. Run JBoss Operations Network (JON) Discovery

Run the `discovery --full` command in the JBoss Operations Network (JON) agent.

### 4. Locate Application Server Process

In the JBoss Operations Network (JON) web interface, the JBoss Enterprise Application Platform process is listed as a JMX server.

### 5. Import the Process Into Inventory

Import the process into the JBoss Operations Network (JON) inventory.

### 6. Optional: Run Discovery Again

If required, run the `discovery --full` command again to discover the new resources.

## Result

The JBoss Data Grid Library mode application is now deployed in JBoss Enterprise Application Platform's standalone mode and can be monitored using the JBoss Operations Network (JON).

[Report a bug](#)

### 22.7.2.3.2. Monitor an Application Deployed in Domain Mode

Use the following instructions to monitor an application deployed in JBoss Enterprise Application Platform 6 using its domain mode:

## Procedure 22.7. Monitor an Application Deployed in Domain Mode

### 1. Edit the Host Configuration

Edit the `domain/configuration/host.xml` file to replace the `server` element with the following configuration:

```
<servers>
  <server name="server-one" group="main-server-group">
    <jvm name="default">
      <jvm-options>
        <option value="-Dorg.rhq.resourceKey=EAP1"/>
      </jvm-options>
    </jvm>
  </server>
  <server name="server-two" group="main-server-group" auto-
start="true">
    <socket-bindings port-offset="150"/>
    <jvm name="default">
      <jvm-options>
        <option value="-Dorg.rhq.resourceKey=EAP2"/>
      </jvm-options>
    </jvm>
  </server>
</servers>
```

### 2. Start JBoss Enterprise Application Platform 6

Start JBoss Enterprise Application Platform 6 in domain mode:

```
$JBOSS_HOME/bin/domain.sh
```

### 3. Deploy the Red Hat JBoss Data Grid Application

Deploy the WAR file that contains the JBoss Data Grid Library mode application with `globalJmxStatistics` and `jmxStatistics` enabled.

### 4. Run Discovery in JBoss Operations Network (JON)

If required, run the `discovery --full` command for the JBoss Operations Network (JON) agent to discover the new resources.

## Result

The JBoss Data Grid Library mode application is now deployed in JBoss Enterprise Application Platform's domain mode and can be monitored using the JBoss Operations Network (JON).

[Report a bug](#)

## 22.8. JBOSS OPERATIONS NETWORK PLUG-IN QUICKSTART

For testing or demonstrative purposes with a single JBoss Operations Network agent, upload the plug-in to the server then type "plugins update" at the agent command line to force a retrieval of the latest plugins from the server.

[Report a bug](#)

## 22.9. OTHER MANAGEMENT TOOLS AND OPERATIONS

Managing Red Hat JBoss Data Grid instances requires exposing significant amounts of relevant statistical information. This information allows administrators to get a clear view of each JBoss Data Grid node's state. A single installation can comprise of tens or hundreds of JBoss Data Grid nodes and it is important to provide this information in a clear and concise manner. JBoss Operations Network is one example of a tool that provides runtime visibility. Other tools, such as **JConsole** can be used where JMX is enabled.

[Report a bug](#)

### 22.9.1. Accessing Data via URLs

Caches that have been configured with a REST interface have access to Red Hat JBoss Data Grid using RESTful HTTP access.

The RESTful service only requires a HTTP client library, eliminating the need for tightly coupled client libraries and bindings. For more information about how to retrieve data using the REST interface, see [Section 11.6, "Using the REST Interface"](#).

HTTP `put()` and `post()` methods place data in the cache, and the URL used determines the cache name and key(s) used. The data is the value placed into the cache, and is placed in the body of the request.

A Content-Type header must be set for these methods. **GET** and **HEAD** methods are used for data retrieval while other headers control cache settings and behavior.

**NOTE**

It is not possible to have conflicting server modules interact with the data grid. Caches must be configured with a compatible interface in order to have access to JBoss Data Grid.

[Report a bug](#)

## 22.9.2. Limitations of Map Methods

Specific Map methods, such as `size()`, `values()`, `keySet()` and `entrySet()`, can be used with certain limitations with Red Hat JBoss Data Grid as they are unreliable. These methods do not acquire locks (global or local) and concurrent modification, additions and removals are excluded from consideration in these calls. Furthermore, the listed methods are only operational on the local cache and do not provide a global view of state.

If the listed methods acted globally, it would result in a significant impact on performance and would produce a scalability bottleneck. As a result, it is recommended that these methods are used for informational and debugging purposes only.

### Performance Concerns

From Red Hat JBoss Data Grid 6.3 onwards, the map methods `size()`, `values()`, `keySet()`, and `entrySet()` include entries in the cache loader by default whereas previously these methods only included the local data container. The underlying cache loader directly affects the performance of these commands. As an example, when using a database, these methods run a complete scan of the table where data is stored which can result in slower processing. Use `Cache.getAdvancedCache().withFlags(Flag.SKIP_CACHE_LOAD).values()` to maintain the old behavior and not loading from the cache loader which would avoid the slower performance.

### Changes to the `size()` Method (Embedded Caches)

In JBoss Data Grid 6.3 the `Cache#size()` method returned only the number of entries on the local node, ignoring other nodes for clustered caches and including any expired entries. While the default behavior was not changed in JBoss Data Grid 6.4 or later, accurate results can be enabled for bulk operations, including `size()`, by setting the `infinispan.accurate.bulk.ops` system property to true. In this mode of operation, the result returned by the `size()` method is affected by the flags `org.infinispan.context.Flag#CACHE_MODE_LOCAL`, to force it to return the number of entries present on the local node, and `org.infinispan.context.Flag#SKIP_CACHE_LOAD`, to ignore any passivated entries.

### Changes to the `size()` Method (Remote Caches)

In JBoss Data Grid 6.3, the Hot Rod `size()` method obtained the size of a cache by invoking the `STATS` operation and using the returned `numberOfEntries` statistic. This statistic is not an accurate measurement of the number of entries in a cache because it does not take into account expired and passivated entries and it is only local to the node that responded to the operation. As an additional result, when security was enabled, the client would need the `ADMIN` permission instead of the more appropriate `BULK_READ`.

In JBoss Data Grid 6.4 and later the Hot Rod protocol has been enhanced with a dedicated `SIZE` operation, and the clients have been updated to use this operation for the `size()` method. The JBoss Data Grid server will need to be started with the `infinispan.accurate.bulk.ops` system property set to `true` so that size can be computed accurately.

[Report a bug](#)



## PART XI. COMMAND LINE TOOLS

Red Hat JBoss Data Grid includes two command line tools for interacting with the caches in the data grid:

- The JBoss Data Grid Library CLI. For more information, see [Section 23.1, “Red Hat JBoss Data Grid Library Mode CLI”](#).
- The JBoss Data Grid Server CLI. For more information, see [Section 23.2, “Red Hat Data Grid Server CLI”](#).

[Report a bug](#)

## CHAPTER 23. RED HAT JBOSS DATA GRID CLIS

Red Hat JBoss Data Grid includes two Command Line Interfaces: a Library Mode CLI (see [Section 23.1, “Red Hat JBoss Data Grid Library Mode CLI”](#) for details) and a Server Mode CLI (see [Section 23.2, “Red Hat Data Grid Server CLI”](#) for details).

[Report a bug](#)

### 23.1. RED HAT JBOSS DATA GRID LIBRARY MODE CLI

Red Hat JBoss Data Grid includes the Red Hat JBoss Data Grid Library Mode Command Line Interface (CLI) that is used to inspect and modify data within caches and internal components (such as transactions, cross-datacenter replication sites, and rolling upgrades). The JBoss Data Grid Library Mode CLI can also be used for more advanced operations such as transactions.

[Report a bug](#)

#### 23.1.1. Start the Library Mode CLI (Server)

Start the Red Hat JBoss Data Grid CLI's server-side module with the **standalone** and **cluster** files. For Linux, use the **standalone.sh** or **clustered.sh** script and for Windows, use the **standalone.bat** or **clustered.bat** file.

[Report a bug](#)

#### 23.1.2. Start the Library Mode CLI (Client)

Start the Red Hat JBoss Data Grid CLI client using the **cli** files in the **bin** directory. For Linux, run **bin/cli.sh** and for Windows, run **bin\cli.bat**.

When starting up the CLI client, specific command line switches can be used to customize the start up.

[Report a bug](#)

#### 23.1.3. CLI Client Switches for the Command Line

The listed command line switches are appended to the command line when starting the Red Hat JBoss Data Grid CLI command:

**Table 23.1. CLI Client Command Line Switches**

Short Option	Long Option	Description
-c	--connect=\${URL}	Connects to a running Red Hat JBoss Data Grid instance. For example, for JMX over RMI use <b>jmx://[username[:password]]@host:port[/container[/cache]]</b> and for JMX over JBoss Remoting use <b>remoting://[username[:password]]@host:port[/container[/cache]]</b>

Short Option	Long Option	Description
-f	--file={FILE}	Read the input from the specified file rather than using interactive mode. If the value is set to - then the <i>stdin</i> is used as the input.
-h	--help	Displays the help information.
-v	--version	Displays the CLI version information.

[Report a bug](#)

### 23.1.4. Connect to the Application

Use the following command to connect to the application using the CLI:

```
[disconnected//]> connect jmx://localhost:12000
[jmx://localhost:12000/MyCacheManager/>
```



#### NOTE

The port value **12000** depends on the value the JVM is started with. For example, starting the JVM with the **-Dcom.sun.management.jmxremote.port=12000** command line parameter uses this port, but otherwise a random port is chosen. When the remoting protocol (**remoting://localhost:9999**) is used, the Red Hat JBoss Data Grid server administration port is used (the default is port **9999**).

The command line prompt displays the active connection information, including the currently selected **CacheManager**.

Use the **cache** command to select a cache before performing cache operations. The CLI supports tab completion, therefore using the **cache** and pressing the tab button displays a list of active caches:

```
[[jmx://localhost:12000/MyCacheManager/> cache
__defaultcache namedCache
[jmx://localhost:12000/MyCacheManager/]> cache __defaultcache
[jmx://localhost:12000/MyCacheManager/__defaultcache]>
```

Additionally, pressing tab displays a list of valid commands for the CLI.

[Report a bug](#)

## 23.2. RED HAT DATA GRID SERVER CLI

Red Hat JBoss Data Grid includes a new Remote Client-Server mode CLI. This CLI can only be used for specific use cases, such as manipulating the server subsystem for the following:

- configuration

- management
- obtaining metrics

[Report a bug](#)

### 23.2.1. Start the Server Mode CLI

Use the following commands to run the JBoss Data Grid Server CLI from the command line:

For Linux:

```
$ JDG_HOME/bin/cli.sh
```

For Windows:

```
C:\>JDG_HOME\bin\cli.bat
```

[Report a bug](#)

## 23.3. CLI COMMANDS

Unless specified otherwise, all listed commands for the JBoss Data Grid CLIs can be used with both the Library Mode and Server Mode CLIs. Specifically, the **deny** (see [Section 23.3.8, “The deny Command”](#)), **grant** (see [Section 23.3.14, “The grant Command”](#)), and **roles** (see [Section 23.3.19, “The roles command”](#)) commands are only available on the Server Mode CLI.

[Report a bug](#)

### 23.3.1. The abort Command

The **abort** command aborts a running batch initiated using the **start** command. Batching must be enabled for the specified cache. The following is a usage example:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> start
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> abort
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
null
```

[Report a bug](#)

### 23.3.2. The begin Command

The **begin** command starts a transaction. This command requires transactions enabled for the cache it targets. An example of this command's usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> begin
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> put b b
[jmx://localhost:12000/MyCacheManager/namedCache]> commit
```

[Report a bug](#)

### 23.3.3. The cache Command

The **cache** command specifies the default cache used for all subsequent operations. If invoked without any parameters, it shows the currently selected cache. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> cache ___defaultcache
[jmx://localhost:12000/MyCacheManager/___defaultcache]> cache
___defaultcache
[jmx://localhost:12000/MyCacheManager/___defaultcache]>
```

[Report a bug](#)

### 23.3.4. The clear Command

The **clear** command clears all content from the cache. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> clear
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
null
```

[Report a bug](#)

### 23.3.5. The commit Command

The **commit** command commits changes to an ongoing transaction. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> begin
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> put b b
[jmx://localhost:12000/MyCacheManager/namedCache]> commit
```

[Report a bug](#)

### 23.3.6. The container Command

The **container** command selects the default cache container (cache manager). When invoked without any parameters, it lists all available containers. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> container
MyCacheManager OtherCacheManager
[jmx://localhost:12000/MyCacheManager/namedCache]> container
OtherCacheManager
[jmx://localhost:12000/OtherCacheManager/]>
```

[Report a bug](#)

### 23.3.7. The create Command

The **create** command creates a new cache based on the configuration of an existing cache definition. An example of its usage is as follows:

■

```
[jmx://localhost:12000/MyCacheManager/namedCache]> create newCache like
namedCache
[jmx://localhost:12000/MyCacheManager/namedCache]> cache newCache
[jmx://localhost:12000/MyCacheManager/newCache]>
```

[Report a bug](#)

### 23.3.8. The deny Command

When authorization is enabled and the role mapper has been configured to be the ClusterRoleMapper, principal to role mappings are stored within the cluster registry (a replicated cache available to all nodes). The **deny** command can be used to deny roles previously assigned to a principal:

```
[remoting://localhost:9999]> deny supervisor to user1
```



#### NOTE

The **deny** command is only available to the JBoss Data Grid Server Mode CLI.

[Report a bug](#)

### 23.3.9. The disconnect Command

The **disconnect** command disconnects the currently active connection, which allows the CLI to connect to another instance. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> disconnect
[disconnected//]
```

[Report a bug](#)

### 23.3.10. The encoding Command

The **encoding** command sets a default codec to use when reading and writing entries to and from a cache. If invoked with no arguments, the currently selected codec is displayed. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> encoding
none
[jmx://localhost:12000/MyCacheManager/namedCache]> encoding --list
memcached
hotrod
none
rest
[jmx://localhost:12000/MyCacheManager/namedCache]> encoding hotrod
```

[Report a bug](#)

### 23.3.11. The end Command

The **end** command ends a running batch initiated using the **start** command. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> start
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> end
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
a
```

[Report a bug](#)

### 23.3.12. The evict Command

The **evict** command evicts an entry associated with a specific key from the cache. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> evict a
```

[Report a bug](#)

### 23.3.13. The get Command

The **get** command shows the value associated with a specified key. For primitive types and Strings, the **get** command prints the default representation. For other objects, a JSON representation of the object is printed. An example of its usage is as follows:


```
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
a
```

[Report a bug](#)

### 23.3.14. The grant Command

When authorization is enabled and the role mapper has been configured to be the **ClusterRoleMapper**, the principal to role mappings are stored within the cluster registry (a replicated cache available to all nodes). The **grant** command can be used to grant new roles to a principal as follows:

```
[remoting://localhost:9999]> grant supervisor to user1
```



#### NOTE

The **grant** command is only available to the JBoss Data Grid Server Mode CLI.

[Report a bug](#)

### 23.3.15. The info Command

The **info** command displays the configuration of a selected cache or container. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> info
```

```
GlobalConfiguration{asyncListenerExecutor=ExecutorFactoryConfiguration{factory=org.infinispan.executors.DefaultExecutorFactory@98add58},
asyncTransportExecutor=ExecutorFactoryConfiguration{factory=org.infinispan
.executors.DefaultExecutorFactory@7bc9c14c},
evictionScheduledExecutor=ScheduledExecutorFactoryConfiguration{factory=or
g.infinispan.executors.DefaultScheduledExecutorFactory@7ab1a411},
replicationQueueScheduledExecutor=ScheduledExecutorFactoryConfiguration{fa
ctory=org.infinispan.executors.DefaultScheduledExecutorFactory@248a9705},
globalJmxStatistics=GlobalJmxStatisticsConfiguration{allowDuplicateDomains
=true, enabled=true, jmxDomain='jboss.infinispan',
mBeanServerLookup=org.jboss.as.clustering.infinispan.MBeanServerProvider@6
c0dc01, cacheManagerName='local', properties={}},
transport=TransportConfiguration{clusterName='ISPN', machineId='null',
rackId='null', siteId='null', strictPeerToPeer=false,
distributedSyncTimeout=240000, transport=null, nodeName='null',
properties={}},
serialization=SerializationConfiguration{advancedExternalizers=
{1100=org.infinispan.server.core.CacheValue$Externalizer@5fab9c91d,
1101=org.infinispan.server.memcached.MemcachedValue$Externalizer@720bffd,
1104=org.infinispan.server.hotrod.ServerAddress$Externalizer@771c7eb2},
marshaller=org.infinispan.marshall.VersionAwareMarshaller@6fc21535,
version=52,
classResolver=org.jboss.marshalling.ModularClassResolver@2efe83e5},
shutdown=ShutdownConfiguration{hookBehavior=DONT_REGISTER}, modules={},
site=SiteConfiguration{localSite='null'}}
```

[Report a bug](#)

### 23.3.16. The locate Command

The **locate** command displays the physical location of a specified entry in a distributed cluster. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> locate a
[host/node1,host/node2]
```

[Report a bug](#)

### 23.3.17. The put Command

The **put** command inserts an entry into the cache. If a mapping exists for a key, the **put** command overwrites the old value. The CLI allows control over the type of data used to store the key and value. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> put b 100
[jmx://localhost:12000/MyCacheManager/namedCache]> put c 41391
[jmx://localhost:12000/MyCacheManager/namedCache]> put d true
[jmx://localhost:12000/MyCacheManager/namedCache]> put e {
"package.MyClass": {"i": 5, "x": null, "b": true } }
```

Optionally, the **put** can specify a life span and maximum idle time value as follows:



```
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a expires 10s
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a expires 10m
maxidle 1m
```

[Report a bug](#)

### 23.3.18. The replace Command

The **replace** command replaces an existing entry in the cache with a specified new value. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> replace a b
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
b
[jmx://localhost:12000/MyCacheManager/namedCache]> replace a b c
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
c
[jmx://localhost:12000/MyCacheManager/namedCache]> replace a b d
[jmx://localhost:12000/MyCacheManager/namedCache]> get a
c
```

[Report a bug](#)

### 23.3.19. The roles command

When authorization is enabled and the role mapper has been configured to be the **ClusterRoleMapper**, the principal to role mappings are stored within the cluster registry (a replicated cache available to all nodes). The **roles** command can be used to list the roles associated to a specific user, or to all users if one is not given:

```
[remoting://localhost:9999]> roles user1
[supervisor, reader]
```



#### NOTE

The **roles** command is only available to the JBoss Data Grid Server Mode CLI.

[Report a bug](#)

### 23.3.20. The rollback Command

The **rollback** command rolls back any changes made by an ongoing transaction. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> begin
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> put b b
[jmx://localhost:12000/MyCacheManager/namedCache]> rollback
```

[Report a bug](#)

### 23.3.21. The site Command

The **site** command performs administration tasks related to cross-datacenter replication. This command also retrieves information about the status of a site and toggles the status of a site. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> site --status NYC
online
[jmx://localhost:12000/MyCacheManager/namedCache]> site --offline NYC
ok
[jmx://localhost:12000/MyCacheManager/namedCache]> site --status NYC
offline
[jmx://localhost:12000/MyCacheManager/namedCache]> site --online NYC
```

[Report a bug](#)

### 23.3.22. The start Command

The **start** command initiates a batch of operations. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> start
[jmx://localhost:12000/MyCacheManager/namedCache]> put a a
[jmx://localhost:12000/MyCacheManager/namedCache]> put b b
[jmx://localhost:12000/MyCacheManager/namedCache]> end
```

[Report a bug](#)

### 23.3.23. The stats Command

The **stats** command displays statistics for the cache. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> stats
Statistics: {
  averageWriteTime: 143
  evictions: 10
  misses: 5
  hitRatio: 1.0
  readWriteRatio: 10.0
  removeMisses: 0
  timeSinceReset: 2123
  statisticsEnabled: true
  stores: 100
  elapsedTime: 93
  averageReadTime: 14
  removeHits: 0
  numberOfEntries: 100
  hits: 1000
}
LockManager: {
  concurrencyLevel: 1000
  numberOfLocksAvailable: 0
  numberOfLocksHeld: 0
}
```

[Report a bug](#)

### 23.3.24. The upgrade Command

The **upgrade** command implements the rolling upgrade procedure. For details about rolling upgrades, see the *Rolling Upgrades* chapter in the Red Hat JBoss Data Grid *Developer Guide*.

An example of the **upgrade** command's use is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> upgrade --  
synchronize=hotrod --all  
[jmx://localhost:12000/MyCacheManager/namedCache]> upgrade --  
disconnectsource=hotrod --all
```

[Report a bug](#)

### 23.3.25. The version Command

The **version** command displays version information for the CLI client and server. An example of its usage is as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> version  
Client Version 5.2.1.Final  
Server Version 5.2.1.Final
```

[Report a bug](#)

## **PART XII. OTHER RED HAT JBOSS DATA GRID FUNCTIONS**

## CHAPTER 24. SET UP THE L1 CACHE

### 24.1. ABOUT THE L1 CACHE

The Level 1 (or L1) cache stores remote cache entries after they are initially accessed, preventing unnecessary remote fetch operations for each subsequent use of the same entries. The L1 cache is only available when Red Hat JBoss Data Grid's cache mode is set to distribution. In other cache modes any configuration related to the L1 cache is ignored.

When caches are configured with distributed mode, the entries are evenly distributed between all clustered caches. Each entry is copied to a desired number of owners, which can be less than the total number of caches. As a result, the system's scalability is improved but also means that some entries are not available on all nodes and must be fetched from their owner node. In this situation, configure the Cache component to use the L1 Cache to temporarily store entries that it does not own to prevent repeated fetching for subsequent uses.

Each time a key is updated an invalidation message is generated. This message is multicast to each node that contains data that corresponds to current L1 cache entries. The invalidation message ensures that each of these nodes marks the relevant entry as invalidated. Also, when the location of an entry changes in the cluster, the corresponding L1 cache entry is invalidated to prevent outdated cache entries.

[Report a bug](#)

### 24.2. L1 CACHE CONFIGURATION

#### 24.2.1. L1 Cache Configuration (Library Mode)

The following sample configuration shows the L1 cache default values in Red Hat JBoss Data Grid's Library Mode.

##### Example 24.1. L1 Cache Configuration in Library Mode

```
<clustering mode="dist">
  <sync/>
  <l1 enabled="true"
      lifespan="60000" />
</clustering>
```

The **l1** element configures the cache behavior in distributed cache instances. If used with non-distributed caches, this element is ignored.

- The **enabled** parameter enables the L1 cache.
- The **lifespan** parameter sets the maximum life span of an entry when it is placed in the L1 cache.

[Report a bug](#)

#### 24.2.2. L1 Cache Configuration (Remote Client-Server Mode)


The following sample configuration shows the L1 cache default values in Red Hat JBoss Data Grid's Remote Client-Server mode.

### Example 24.2. L1 Cache Configuration for Remote Client-Server Mode

```
<distributed-cache l1-lifespan="${VALUE}">  
  <!-- Additional configuration information here -->  
</distributed-cache>
```

The **l1-lifespan** element is added to a **distributed-cache** element to enable L1 caching and to set the life span of the L1 cache entries for the cache. This element is only valid for distributed caches.

If **l1-lifespan** is set to **0** or a negative number (**-1**), L1 caching is disabled. L1 caching is enabled when the **l1-lifespan** value is greater than **0**.



#### NOTE

When the cache is accessed remotely via the Hot Rod protocol, the client accesses the owner node directly. Therefore, using L1 Cache in this situation does not offer any performance improvement and is not recommended. Other remote clients (Memcached, REST) cannot target the owner, therefore, using L1 Cache may increase the performance (at the cost of higher memory consumption).



#### NOTE

In Remote Client-Server mode, the L1 cache was enabled by default when distributed cache was used, even if the **l1-lifespan** attribute is not set. The default lifespan value was 10 minutes. Since JBoss Data Grid 6.3, the default lifespan is 0 which disables the L1 cache. Set a non-zero value for the **l1-lifespan** parameter to enable the L1 cache.

[Report a bug](#)

## CHAPTER 25. SET UP TRANSACTIONS

### 25.1. ABOUT TRANSACTIONS

A transaction consists of a collection of interdependent or related operations or tasks. All operations within a single transaction must succeed for the overall success of the transaction. If any operations within a transaction fail, the transaction as a whole fails and rolls back any changes. Transactions are particularly useful when dealing with a series of changes as part of a larger operation.

In Red Hat JBoss Data Grid, transactions are only available in Library mode.

[Report a bug](#)

#### 25.1.1. About the Transaction Manager

In Red Hat JBoss Data Grid, the Transaction Manager coordinates transactions across a single or multiple resources. The responsibilities of a Transaction Manager include:

- initiating and concluding transactions
- managing information about each transaction
- coordinating transactions as they operate over multiple resources
- recovering from a failed transaction by rolling back changes

[Report a bug](#)

#### 25.1.2. XA Resources and Synchronizations

XA Resources are fully fledged transaction participants. In the prepare phase (see [Section F.7, “About Two Phase Commit \(2PC\)”](#) for details), the XA Resource returns a vote with either the value **OK** or **ABORT**. If the Transaction Manager receives **OK** votes from all XA Resources, the transaction is committed, otherwise it is rolled back.

Synchronizations are a type of listener that receive notifications about events leading to the transaction life cycle. Synchronizations receive an event before and after the operation completes.

Unless recovery is required, it is not necessary to register as a full XA resource. An advantage of synchronizations is that they allow the Transaction Manager to optimize 2PC (Two Phase Commit) with a 1PC (One Phase Commit) where only one other resource is enlisted with that transaction (last resource commit optimization). This makes registering a synchronization more efficient.

However, if the operation fails in the prepare phase within Red Hat JBoss Data Grid, the transaction is not rolled back and if there are more participants in the transaction, they can ignore this failure and commit. Additionally, errors encountered in the commit phase are not propagated to the application code that commits the transaction.

By default JBoss Data Grid registers to the transaction as a synchronization.

[Report a bug](#)

#### 25.1.3. Optimistic and Pessimistic Transactions

Pessimistic transactions acquire the locks when the first write operation on the key executes. After the key is locked, no other transaction can modify the key until this transaction is committed or rolled back. It is up to the application code to acquire the locks in correct order to prevent deadlocks.

With optimistic transactions locks are acquired at transaction prepare time and are held until the transaction commits (or rolls back). Also, Red Hat JBoss Data Grid sorts keys for all entries modified within a transaction automatically, preventing any deadlocks occurring due to the incorrect order of keys being locked. This results in:

- less messages being sent during the transaction execution
- locks held for shorter periods
- improved throughput



#### NOTE

Read operations never acquire any locks. Acquiring the lock for a read operation on demand is possible only with pessimistic transactions, using the **FORCE\_WRITE\_LOCK** flag with the operation.

[Report a bug](#)

### 25.1.4. Write Skew Checks

A common use case for entries is that they are read and subsequently written in a transaction. However, a third transaction can modify the entry between these two operations. In order to detect such a situation and roll back a transaction Red Hat JBoss Data Grid offers entry versioning and write skew checks. If the modified version is not the same as when it was last read during the transaction, the write skew checks throws an exception and the transaction is rolled back.

Enabling write skew check requires the **REPEATABLE\_READ** isolation level. Also, in clustered mode (distributed or replicated modes), set up entry versioning. For local mode, entry versioning is not required.



#### IMPORTANT

With optimistic transactions, write skew checks are required for (atomic) conditional operations.

[Report a bug](#)

### 25.1.5. Transactions Spanning Multiple Cache Instances

Each cache operates as a separate, standalone Java Transaction API (JTA) resource. However, components can be internally shared by Red Hat JBoss Data Grid for optimization, but this sharing does not affect how caches interact with a Java Transaction API (JTA) Manager.

[Report a bug](#)

## 25.2. CONFIGURE TRANSACTIONS



### 25.2.1. Configure Transactions (Library Mode)

In Red Hat JBoss Data Grid, transactions in Library mode can be configured with synchronization and transaction recovery. Transactions in their entirety (which includes synchronization and transaction recovery) are not available in Remote Client-Server mode.

In Library mode, transactions are configured as follows:

#### Procedure 25.1. Configure Transactions in Library Mode (XML Configuration)

```
<namedCache <!-- Additional configuration information here -->>
  <transaction <!-- Additional configuration information here --> >
  <locking <!-- Additional configuration information here --> >
  <versioning enabled="{true,false}"
    versioningScheme="{NONE|SIMPLE}"/>
    <!-- Additional configuration information here -->
</namedCache>
```

1. Set the **versioning** parameter's **enabled** parameter to **true**.
2. Set the **versioningScheme** parameter to either **NONE** or **SIMPLE** to set the versioning scheme used.

#### Procedure 25.2. Configure Transactions in Library Mode (Programmatic Configuration)

```
1. Configuration config = new ConfigurationBuilder()/* ...
   */.transaction()
      .transactionMode(TransactionMode.TRANSACTIONAL)
      .transactionManagerLookup(new
GenericTransactionManagerLookup())
      .lockingMode(LockingMode.OPTIMISTIC)
      .useSynchronization(true)
      .recovery()

      .recoveryInfoCacheName("anotherRecoveryCacheName").build();
```

- a. Set the transaction mode.
- b. Select and set a lookup class. See the table below this procedure for a list of available lookup classes.
- c. The **lockingMode** value determines whether optimistic or pessimistic locking is used. If the cache is non-transactional, the locking mode is ignored. The default value is **OPTIMISTIC**.
- d. The **useSynchronization** value configures the cache to register a synchronization with the transaction manager, or register itself as an XA resource. The default value is **true** (use synchronization).
- e. The **recovery** parameter enables recovery for the cache when set to **true**.

The **recoveryInfoCacheName** sets the name of the cache where recovery information is held. The default name of the cache is specified by **RecoveryConfiguration.DEFAULT\_RECOVERY\_INFO\_CACHE**.

#### 2. Configure Write Skew Check

The *writeSkew* check determines if a modification to the entry from a different transaction should roll back the transaction. Write skew set to **true** requires *isolation\_level* set to **REPEATABLE\_READ**. The default value for *writeSkew* and *isolation\_level* are **false** and **READ\_COMMITTED** respectively.

```
Configuration config = new ConfigurationBuilder()/* ... */.locking()
    .isolationLevel(IsolationLevel.REPEATABLE_READ).writeSkewCheck(true)
    ;
```

### 3. Configure Entry Versioning

For clustered caches, enable write skew check by enabling entry versioning and setting its value to **SIMPLE**.

```
Configuration config = new ConfigurationBuilder()/* ...
*/.versioning()
    .enable()
    .scheme(VersioningScheme.SIMPLE);
```

**Table 25.1. Transaction Manager Lookup Classes**

Class Name	Details
org.infinispan.transaction.lookup.DummyTransactionManagerLookup	Used primarily for testing environments. This testing transaction manager is not for use in a production environment and is severely limited in terms of functionality, specifically for concurrent transactions and recovery.
org.infinispan.transaction.lookup.JBossStandaloneJTAManagerLookup	The default transaction manager when Red Hat JBoss Data Grid runs in a standalone environment. It is a fully functional JBoss Transactions based transaction manager that overcomes the functionality limits of the <b>DummyTransactionManager</b> .
org.infinispan.transaction.lookup.GenericTransactionManagerLookup	GenericTransactionManagerLookup is used by default when no transaction lookup class is specified. This lookup class is recommended when using JBoss Data Grid with Java EE-compatible environment that provides a TransactionManager interface, and is capable of locating the Transaction Manager in most Java EE application servers. If no transaction manager is located, it defaults to <b>DummyTransactionManager</b> .
org.infinispan.transaction.lookup.JBossTransactionManagerLookup	The <b>JbossTransactionManagerLookup</b> finds the standard transaction manager running in the application server. This lookup class uses JNDI to look up the TransactionManager instance, and is recommended when custom caches are being used in JTA transactions.

[Report a bug](#)

## 25.2.2. Configure Transactions (Remote Client-Server Mode)


Red Hat JBoss Data Grid does not offer transactions in Remote Client-Server mode. The default and only supported configuration is non-transactional, which is set as follows:

### Example 25.1. Transaction Configuration in Remote Client-Server Mode

```

<cache>
  <!-- Additional configuration elements here -->
  <transaction mode="NONE" />
  <!-- Additional configuration elements here -->
</cache>

```



### IMPORTANT

In Remote Client-Server mode, the transactions element is set to **NONE** unless JBoss Data Grid is used in a compatibility mode where the cluster contains both JBoss Data Grid server and library instances. In this case, if transactions are configured in the library mode instance, they must also be configured in the server instance.

[Report a bug](#)

## 25.3. TRANSACTION RECOVERY

The Transaction Manager coordinates the recovery process and works with Red Hat JBoss Data Grid to determine which transactions require manual intervention to complete operations. This process is known as transaction recovery.

JBoss Data Grid uses JMX for operations that explicitly force transactions to commit or roll back. These methods receive byte arrays that describe the XID instead of the number associated with the relevant transactions.

The System Administrator can use such JMX operations to facilitate automatic job completion for transactions that require manual intervention. This process uses the Transaction Manager's transaction recovery process and has access to the Transaction Manager's XID objects.

[Report a bug](#)

### 25.3.1. Transaction Recovery Process

The following process outlines the transaction recovery process in Red Hat JBoss Data Grid.

#### Procedure 25.3. The Transaction Recovery Process

1. The Transaction Manager creates a list of transactions that require intervention.
2. The system administrator, connected to JBoss Data Grid using JMX, is presented with the list of transactions (including transaction IDs) using email or logs. The status of each transaction is either **COMMITTED** or **PREPARED**. If some transactions are in both **COMMITTED** and **PREPARED**

states, it indicates that the transaction was committed on some nodes while in the preparation state on others.

3. The System Administrator visually maps the XID received from the Transaction Manager to a JBoss Data Grid internal ID. This step is necessary because the XID (a byte array) cannot be conveniently passed to the JMX tool and then reassembled by JBoss Data Grid without this mapping.
4. The system administrator forces the commit or rollback process for a transaction based on the mapped internal ID.

[Report a bug](#)

### 25.3.2. Transaction Recovery Example

The following example describes how transactions are used in a situation where money is transferred from an account stored in a database to an account stored in Red Hat JBoss Data Grid.

#### Example 25.2. Money Transfer from an Account Stored in a Database to an Account in JBoss Data Grid

1. The `TransactionManager.commit()` method is invoked to run the two phase commit protocol between the source (the database) and the destination (JBoss Data Grid) resources.
2. The `TransactionManager` tells the database and JBoss Data Grid to initiate the prepare phase (the first phase of a Two Phase Commit).

During the commit phase, the database applies the changes but JBoss Data Grid fails before receiving the Transaction Manager's commit request. As a result, the system is in an inconsistent state due to an incomplete transaction. Specifically, the amount to be transferred has been subtracted from the database but is not yet visible in JBoss Data Grid because the prepared changes could not be applied.

Transaction recovery is used here to reconcile the inconsistency between the database and JBoss Data Grid entries.



#### NOTE

To use JMX to manage transaction recoveries, JMX support must be explicitly enabled.

[Report a bug](#)

## 25.4. DEADLOCK DETECTION

A deadlock occurs when multiple processes or tasks wait for the other to release a mutually required resource. Deadlocks can significantly reduce the throughput of a system, particularly when multiple transactions operate against one key set.

Red Hat JBoss Data Grid provides deadlock detection to identify such deadlocks. Deadlock detection is set to **disabled** by default.

[Report a bug](#)

### 25.4.1. Enable Deadlock Detection

Deadlock detection in Red Hat JBoss Data Grid is set to **disabled** by default but can be enabled and configured for each cache using the *namedCache* configuration element by adding the following:

```
<deadlockDetection enabled="true" spinDuration="100"/>
```

The *spinDuration* attribute defines how often lock acquisition is attempted within the maximum time allowed to acquire a particular lock (in milliseconds).

Deadlock detection can only be applied to individual caches. Deadlocks that are applied on more than one cache cannot be detected by JBoss Data Grid.

[Report a bug](#)

## CHAPTER 26. CONFIGURE JGROUPS

JGroups is the underlying group communication library used to connect Red Hat JBoss Data Grid instances. For a full list of JGroups protocols supported in JBoss Data Grid, see [Section A.1, “Supported JGroups Protocols”](#)

[Report a bug](#)

### 26.1. CONFIGURE RED HAT JOSS DATA GRID INTERFACE BINDING (REMOTE CLIENT-SERVER MODE)

#### 26.1.1. Interfaces

Red Hat JBoss Data Grid allows users to specify an interface type rather than a specific (unknown) IP address.

- **link-local**: Uses a **169.x.x.x** or **254.x.x.x** address. This suits the traffic within one box.

```
<interfaces>
  <interface name="link-local">
    <link-local-address/>
  </interface>
  <!-- Additional configuration elements here -->
</interfaces>
```

- **site-local**: Uses a private IP address, for example **192.168.x.x**. This prevents extra bandwidth charged from GoGrid, and similar providers.

```
<interfaces>
  <interface name="site-local">
    <site-local-address/>
  </interface>
  <!-- Additional configuration elements here -->
</interfaces>
```

- **global**: Picks a public IP address. This should be avoided for replication traffic.

```
<interfaces>
  <interface name="global">
    <any-address/>
  </interface>
  <!-- Additional configuration elements here -->
</interfaces>
```

- **non-loopback**: Uses the first address found on an active interface that is not a **127.x.x.x** address.

```
<interfaces>
  <interface name="non-loopback">
    <not>
      <loopback />
    </not>
  </interface>
</interfaces>
```

```

    </not>
  </interface>
</interfaces>

```

[Report a bug](#)

## 26.1.2. Binding Sockets

Socket bindings provide a named the combination of interface and port. Sockets can be bound to the interface either individually or using a socket binding group.

[Report a bug](#)

### 26.1.2.1. Binding a Single Socket Example

The following is an example depicting the use of JGroups interface socket binding to bind an individual socket using the *socket-binding* element.

#### Example 26.1. Socket Binding

```

<socket-binding name="jgroups-udp" <!-- Additional configuration
elements here --> interface="site-local"/>

```

[Report a bug](#)

### 26.1.2.2. Binding a Group of Sockets Example

The following is an example depicting the use of Groups interface socket bindings to bind a group, using the *socket-binding-group* element:

#### Example 26.2. Bind a Group

```

<socket-binding-group name="ha-sockets" default-interface="global">
  <!-- Additional configuration elements here -->
  <socket-binding name="jgroups-tcp" port="7600"/>
  <socket-binding name="jgroups-tcp-fd" port="57600"/>
  <!-- Additional configuration elements here -->
</socket-binding-group>

```

The two sample socket bindings in the example are bound to the same *default-interface* (**global**), therefore the interface attribute does not need to be specified.

[Report a bug](#)

## 26.1.3. Configure JGroups Socket Binding

Each JGroups stack, configured in the JGroups subsystem, uses a specific socket binding. Set up the socket binding as follows:

#### Example 26.3. JGroups UDP Socket Binding Configuration

The following example uses UDP to automatically detect additional nodes on the network:

```
<subsystem xmlns="urn:jboss:domain:jgroups:1.2" default-stack="udp">
  <stack name="udp">
    <transport type="UDP" socket-binding="jgroups-udp">
      <!-- Additional configuration elements here -->
    </transport>
    <!-- rest of protocols -->
  </stack>
</subsystem>
```

#### Example 26.4. JGroups TCP Socket Binding Configuration

The following example uses TCP to establish direct communication between two clusters nodes. In the example below node1 is located at 192.168.1.2:7600, and node2 is located at 192.168.1.3:7600. The port in use will be defined by the `jgroups-tcp` property in the **socket-binding** section.

```
<subsystem xmlns="urn:infinispan:server:jgroups:6.1" default-
stack="tcp">
  <stack name="tcp">
    <transport type="TCP" socket-binding="jgroups-tcp"/>
    <protocol type="TCPPING">
      <property
name="initial_hosts">192.168.1.2[7600],192.168.1.3[7600]</property>
      <property name="num_initial_members">2</property>
      <property name="port_range">0</property>
      <property name="timeout">2000</property>
    </protocol>
    <protocol type="MERGE3"/>
    <protocol type="FD SOCK" socket-binding="jgroups-tcp-fd"/>
    <protocol type="FD_ALL"/>
    <protocol type="VERIFY_SUSPECT"/>
    <protocol type="pbcast.NAKACK2">
      <property name="use_mcast_xmit">>false</property>
    </protocol>
    <protocol type="UNICAST3"/>
    <protocol type="pbcast.STABLE"/>
    <protocol type="pbcast.GMS"/>
    <protocol type="MFC"/>
    <protocol type="FRAG2"/>
  </stack>
</subsystem>
```

The decision of UDP vs TCP must be made in each environment. By default JGroups uses UDP, as it allows for dynamic detection of clustered members and scales better in larger clusters due to a smaller network footprint. In addition, when using UDP only one packet per cluster is required, as multicast packets are received by all subscribers to the multicast address; however, in environments where multicast traffic is prohibited, or if UDP traffic can not reach the remote cluster nodes, such as when cluster members are on separate VLANs, TCP traffic can be used to create a cluster.





## IMPORTANT

When using UDP as the JGroups transport, the socket binding has to specify the regular (unicast) port, multicast address, and multicast port.

[Report a bug](#)

## 26.2. CONFIGURE JGROUPS (LIBRARY MODE)

Red Hat JBoss Data Grid must have an appropriate JGroups configuration in order to operate in clustered mode.

### Example 26.5. JGroups Programmatic Configuration

```
GlobalConfiguration gc = new GlobalConfigurationBuilder()
    .transport()
    .defaultTransport()
    .addProperty("configurationFile", "jgroups.xml")
    .build();
```

### Example 26.6. JGroups XML Configuration

```
<infinispan>
  <global>
    <transport>
      <properties>
        <property name="configurationFile" value="jgroups.xml" />
      </properties>
    </transport>
  </global>
  <!-- Additional configuration elements here -->
</infinispan>
```

In either programmatic or XML configuration methods, JBoss Data Grid searches for **jgroups.xml** in the classpath before searching for an absolute path name if it is not found in the classpath.

[Report a bug](#)

### 26.2.1. JGroups Transport Protocols

A transport protocol is the protocol at the bottom of a protocol stack. Transport Protocols are responsible for sending and receiving messages from the network.

Red Hat JBoss Data Grid ships with both UDP and TCP transport protocols.

[Report a bug](#)

#### 26.2.1.1. The UDP Transport Protocol

UDP is a transport protocol that uses:

- IP multicasting to send messages to all members of a cluster.
- UDP datagrams for unicast messages, which are sent to a single member.

When the UDP transport is started, it opens a unicast socket and a multicast socket. The unicast socket is used to send and receive unicast messages, the multicast socket sends and receives multicast sockets. The physical address of the channel will be the same as the address and port number of the unicast socket.

[Report a bug](#)

### 26.2.1.2. The TCP Transport Protocol

TCP/IP is a replacement transport for UDP in situations where IP multicast cannot be used, such as operations over a WAN where routers may discard IP multicast packets.

TCP is a transport protocol used to send unicast and multicast messages.

- When sending multicast messages, TCP sends multiple unicast messages.
- When using TCP, each message to all cluster members is sent as multiple unicast messages, or one to each member.

As IP multicasting cannot be used to discover initial members, another mechanism must be used to find initial membership.

Red Hat JBoss Data Grid's Hot Rod is a custom TCP client/server protocol.

[Report a bug](#)

### 26.2.1.3. Using the TCPPing Protocol

Some networks only allow TCP to be used. The pre-configured **default-configs/default-jgroups-tcp.xml** includes the **MPING** protocol, which uses **UDP** multicast for discovery. When **UDP** multicast is not available, the **MPING** protocol, has to be replaced by a different mechanism. The recommended alternative is the **TCPPING** protocol. The **TCPPING** configuration contains a static list of IP addresses which are contacted for node discovery.

#### Example 26.7. Configure the JGroups Subsystem to Use TCPPING

```
<TCP bind_port="7800" />
<TCPPING
  initial_hosts="${jgroups.tcpping.initial_hosts:HostA[7800],HostB[7801]}"
  port_range="1" />
```

[Report a bug](#)

## 26.2.2. Pre-Configured JGroups Files

Red Hat JBoss Data Grid ships with a number of pre-configured JGroups files packaged in **infinispan-embedded.jar**, and are available on the classpath by default. In order to use one of these files, specify one of these file names instead of using **jgroups.xml**.

The JGroups configuration files shipped with JBoss Data Grid are intended to be used as a starting point for a working project. JGroups will usually require fine-tuning for optimal network performance.

The available configurations are:

- `default-configs/default-jgroups-udp.xml`
- `default-configs/default-jgroups-tcp.xml`
- `default-configs/default-jgroups-ec2.xml`

[Report a bug](#)

### 26.2.2.1. default-jgroups-udp.xml

The `default-configs/default-jgroups-udp.xml` file is a pre-configured JGroups configuration in Red Hat JBoss Data Grid. The `default-jgroups-udp.xml` configuration

- uses UDP as a transport and UDP multicast for discovery.
- is suitable for large clusters (over 8 nodes).
- is suitable if using Invalidation or Replication modes.

The behavior of some of these settings can be altered by adding certain system properties to the JVM at startup. The settings that can be configured are shown in the following table.

**Table 26.1. default-jgroups-udp.xml System Properties**

System Property	Description	Default	Required?
<code>jgroups.udp.mcast_addr</code>	IP address to use for multicast (both for communications and discovery). Must be a valid Class D IP address, suitable for IP multicast.	228.6.7.8	No
<code>jgroups.udp.mcast_port</code>	Port to use for multicast socket	46655	No
<code>jgroups.udp.ip_ttl</code>	Specifies the time-to-live (TTL) for IP multicast packets. The value here refers to the number of network hops a packet is allowed to make before it is dropped	2	No

[Report a bug](#)

### 26.2.2.2. default-jgroups-tcp.xml

The `default-configs/default-jgroups-tcp.xml` file is a pre-configured JGroups configuration in Red Hat JBoss Data Grid. The `default-jgroups-tcp.xml` configuration

- uses TCP as a transport and UDP multicast for discovery.

- is generally only used where multicast UDP is not an option.
- TCP does not perform as well as UDP for clusters of eight or more nodes. Clusters of four nodes or fewer result in roughly the same level of performance for both UDP and TCP.

As with other pre-configured JGroups files, the behavior of some of these settings can be altered by adding certain system properties to the JVM at startup. The settings that can be configured are shown in the following table.

**Table 26.2. default-jgroups-tcp.xml System Properties**

System Property	Description	Default	Required?
jgroups.tcp.address	IP address to use for the TCP transport.	127.0.0.1	No
jgroups.tcp.port	Port to use for TCP socket	7800	No
jgroups.udp.mcast_addr	IP address to use for multicast (for discovery). Must be a valid Class D IP address, suitable for IP multicast.	228.6.7.8	No
jgroups.udp.mcast_port	Port to use for multicast socket	46655	No
jgroups.udp.ip_ttl	Specifies the time-to-live (TTL) for IP multicast packets. The value here refers to the number of network hops a packet is allowed to make before it is dropped	2	No

[Report a bug](#)

### 26.2.2.3. default-jgroups-ec2.xml

The `default-configs/default-jgroups-ec2.xml` file is a pre-configured JGroups configuration in Red Hat JBoss Data Grid. The `default-jgroups-ec2.xml` configuration

- uses TCP as a transport and S3\_PING for discovery.
- is suitable on Amazon EC2 nodes where UDP multicast isn't available.

As with other pre-configured JGroups files, the behavior of some of these settings can be altered by adding certain system properties to the JVM at startup. The settings that can be configured are shown in the following table.

**Table 26.3. default-jgroups-ec2.xml System Properties**

System Property	Description	Default	Required?
jgroups.tcp.address	IP address to use for the TCP transport.	127.0.0.1	No

System Property	Description	Default	Required?
<code>jgroups.tcp.port</code>	Port to use for TCP socket	7800	No
<code>jgroups.s3.access_key</code>	The Amazon S3 access key used to access an S3 bucket		Yes
<code>jgroups.s3.secret_access_key</code>	The Amazon S3 secret key used to access an S3 bucket		Yes
<code>jgroups.s3.bucket</code>	Name of the Amazon S3 bucket to use. Must be unique and must already exist		Yes
<code>jgroups.s3.pre_signed_delete_url</code>	The pre-signed URL to be used for the DELETE operation.		Yes
<code>jgroups.s3.pre_signed_put_url</code>	The pre-signed URL to be used for the PUT operation.		Yes
<code>jgroups.s3.prefix</code>	If set, S3_PING searches for a bucket with a name that starts with the prefix value.		No

[Report a bug](#)

## 26.3. TEST MULTICAST USING JGROUPS

Learn how to ensure that the system has correctly configured multicasting within the cluster.

[Report a bug](#)

### 26.3.1. Testing With Different Red Hat JBoss Data Grid Versions

The following table details which Red Hat JBoss Data Grid versions are compatible with this multicast test:

**Table 26.4. Testing with Different JBoss Data Grid Versions**

Version	Test Case	Details
JBoss Data Grid 6.0.0	Not Available	This version of JBoss Data Grid is based on JBoss Enterprise Application Server 6.0, which does not include the test classes used for this test.

Version	Test Case	Details
JBoss Data Grid 6.0.1	Not Available	This version of JBoss Data Grid is based on JBoss Enterprise Application Platform 6.0, which does not include the test classes used for this test.
JBoss Data Grid 6.1.0	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• In Library mode, they are in the JGroups JAR file in the <b>lib</b> directory.</li> <li>• In Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/org/jgroups/main/</code></b> directory.</li> </ul>
JBoss Data Grid 6.2.0	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• In Library mode, they are in the JGroups JAR file in the <b>lib</b> directory.</li> <li>• In Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/system/layers/base/org/jgroups/main.</code></b></li> </ul>
JBoss Data Grid 6.2.1	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• In Library mode, they are in the JGroups JAR file in the <b>lib</b> directory.</li> <li>• In Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/system/layers/base/org/jgroups/main</code></b></li> </ul>

Version	Test Case	Details
JBoss Data Grid 6.3.0	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• In Library mode, they are in the JGroups JAR file in the lib directory.</li> <li>• In Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/system/layers/base/org/jgroups/main</code></b>.</li> </ul>
JBoss Data Grid 6.4.0	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• For library mode, they are inside the <b><code>infinispan-embedded</code></b> JAR file</li> <li>• For Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/system/layers/base/org/jgroups/main/</code></b> directory."</li> </ul>
JBoss Data Grid 6.5.0	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• For library mode, they are inside the <b><code>infinispan-embedded</code></b> JAR file</li> <li>• For Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/system/layers/base/org/jgroups/main/</code></b> directory."</li> </ul>

Version	Test Case	Details
JBoss Data Grid 6.5.1	Available	<p>The location of the test classes depends on the distribution:</p> <ul style="list-style-type: none"> <li>• For library mode, they are inside the <b>infinispan-embedded</b> JAR file</li> <li>• For Remote Client-Server mode, they are in the JGroups JAR file in the <b><code>\${JDG_HOME}/modules/system/layers/base/org/jgroups/main/</code></b> directory."</li> </ul>

[Report a bug](#)

## 26.3.2. Testing Multicast Using JGroups

The following procedure details the steps to test multicast using JGroups if you are using Red Hat JBoss Data Grid :

### Prerequisites

Ensure that the following prerequisites are met before starting the testing procedure.

1. Set the **`bind_addr`** value to the appropriate IP address for the instance.
2. For added accuracy, set **`mcast_addr`** and **`port`** values that are the same as the cluster communication values.
3. Start two command line terminal windows. Navigate to the location of the JGroups JAR file for one of the two nodes in the first terminal and the same location for the second node in the second terminal.

### Procedure 26.1. Test Multicast Using JGroups

#### 1. Run the Multicast Server on Node One

Run the following command on the command line terminal for the first node (replace **`jgroups.jar`** with the **`infinispan-embedded.jar`** for Library mode):

```
java -cp jgroups.jar org.jgroups.tests.McastReceiverTest -mcast_addr
230.1.2.3 -port 5555 -bind_addr $YOUR_BIND_ADDRESS
```

#### 2. Run the Multicast Server on Node Two

Run the following command on the command line terminal for the second node (replace **`jgroups.jar`** with the **`infinispan-embedded.jar`** for Library mode):

```
java -cp jgroups.jar org.jgroups.tests.McastSenderTest -mcast_addr
230.1.2.3 -port 5555 -bind_addr $YOUR_BIND_ADDRESS
```



**3. Transmit Information Packets**

Enter information on instance for node two (the node sending packets) and press enter to send the information.

**4. View Receives Information Packets**

View the information received on the node one instance. The information entered in the previous step should appear here.

**5. Confirm Information Transfer**

Repeat steps 3 and 4 to confirm all transmitted information is received without dropped packets.

**6. Repeat Test for Other Instances**

Repeat steps 1 to 4 for each combination of sender and receiver. Repeating the test identifies other instances that are incorrectly configured.

**Result**

All information packets transmitted from the sender node must appear on the receiver node. If the sent information does not appear as expected, multicast is incorrectly configured in the operating system or the network.

[Report a bug](#)

## CHAPTER 27. USE RED HAT DATA GRID WITH AMAZON WEB SERVICES

### 27.1. THE S3\_PING JGROUPS DISCOVERY PROTOCOL

**S3\_PING** is a discovery protocol that is ideal for use with Amazon's Elastic Compute Cloud (EC2) because EC2 does not allow multicast and therefore **MPING** is not allowed.

Each EC2 instance adds a small file to an S3 data container, known as a bucket. Each instance then reads the files in the bucket to discover the other members of the cluster.

[Report a bug](#)

### 27.2. S3\_PING CONFIGURATION OPTIONS

Red Hat JBoss Data Grid works with Amazon Web Services in two ways:

- In Library mode, use JGroups' **default-configs/default-jgroups-ec2.xml** file (see [Section 26.2.2.3, "default-jgroups-ec2.xml"](#) for details) or use the **S3\_PING** protocol.
- In Remote Client-Server mode, use JGroups' **S3\_PING** protocol.

In Library and Remote Client-Server mode, there are three ways to configure the **S3\_PING** protocol for clustering to work in Amazon AWS:

- Use Private S3 Buckets. These buckets use Amazon AWS credentials.
- Use Pre-Signed URLs. These pre-assigned URLs are assigned to buckets with private write and public read rights.
- Use Public S3 Buckets. These buckets do not have any credentials.

[Report a bug](#)

#### 27.2.1. Using Private S3 Buckets

This configuration requires access to a private bucket that can only be accessed with the appropriate AWS credentials. To confirm that the appropriate permissions are available, confirm that the user has the following permissions for the bucket:

- List
- Upload/Delete
- View Permissions
- Edit Permissions

Ensure that the **S3\_PING** configuration includes the following properties:

- either the **location** or the **prefix** property to specify the bucket, but not both. If the **prefix** property is set, **S3\_PING** searches for a bucket with a name that starts with the prefix value. If a bucket with the prefix at the beginning of the name is found, **S3\_PING** uses that bucket. If a

bucket with the prefix is not found, **S3\_PING** creates a bucket using the AWS credentials and names it based on the prefix and a UUID (the naming format is *{prefix value}-{UUID}*).

- the **access\_key** and **secret\_access\_key** properties for the AWS user.



## NOTE

If a **403** error displays when using this configuration, verify that the properties have the correct values. If the problem persists, confirm that the system time in the EC2 node is correct. Amazon S3 rejects requests with a time stamp that is more than **15** minutes old compared to their server's times for security purposes.

### Example 27.1. Start the Red Hat JBoss Data Grid Server with a Private Bucket

Run the following command from the top level of the server directory to start the Red Hat JBoss Data Grid server using a private S3 bucket:

```
bin/clustered.sh -Djboss.bind.address={server_ip_address} -
Djboss.bind.address.management={server_ip_address} -
Djboss.default.jgroups.stack=s3 -Djgroups.s3.bucket={s3_bucket_name} -
Djgroups.s3.access_key={access_key} -
Djgroups.s3.secret_access_key={secret_access_key}
```

1. Replace *{server\_ip\_address}* with the server's IP address.
2. Replace *{s3\_bucket\_name}* with the appropriate bucket name.
3. Replace *{access\_key}* with the user's access key.
4. Replace *{secret\_access\_key}* with the user's secret access key.

[Report a bug](#)

## 27.2.2. Using Pre-Signed URLs

For this configuration, create a publically readable bucket in S3 by setting the **List** permissions to **Everyone** to allow public read access. Each node in the cluster generates a pre-signed URL for put and delete operations, as required by the **S3\_PING** protocol. This URL points to a unique file and can include a folder path within the bucket.



## NOTE

Longer paths will cause errors in **S3\_PING**. For example, a path such as **my\_bucket/DemoCluster/node1** works while a longer path such as **my\_bucket/Demo/Cluster/node1** will not.

[Report a bug](#)

### 27.2.2.1. Generating Pre-Signed URLs

JGroup's **S3\_PING** class includes a utility method to generate pre-signed URLs. The last argument for this method is the time when the URL expires expressed in the number of seconds since the Unix epoch (January 1, 1970).

The syntax to generate a pre-signed URL is as follows:

```
String Url = S3_PING.generatePreSignedUrl("{access_key}",
    "{secret_access_key}", "{operation}", "{bucket_name}", "{path}",
    {seconds});
```

1. Replace *{operation}* with either **PUT** or **DELETE**.
2. Replace *{access\_key}* with the user's access key.
3. Replace *{secret\_access\_key}* with the user's secret access key.
4. Replace *{bucket\_name}* with the name of the bucket.
5. Replace *{path}* with the desired path to the file within the bucket.
6. Replace *{seconds}* with the number of seconds since the Unix epoch (January 1, 1970) that the path remains valid.

#### Example 27.2. Generate a Pre-Signed URL

```
String putUrl = S3_PING.generatePreSignedUrl("access_key",
    "secret_access_key", "put", "my_bucket", "DemoCluster/node1",
    1234567890);
```

Ensure that the **S3\_PING** configuration includes the *pre\_signed\_put\_url* and *pre\_signed\_delete\_url* properties generated by the call to **S3\_PING.generatePreSignedUrl()**. This configuration is more secure than one using private S3 buckets, because the AWS credentials are not stored on each node in the cluster



#### NOTE

If a pre-signed URL is entered into an XML file, then the **&** characters in the URL must be replaced with its XML entity (**&amp;**).

[Report a bug](#)

#### 27.2.2.2. Set Pre-Signed URLs Using the Command Line

To set the pre-signed URLs using the command line, use the following guidelines:

- Enclose the URL in double quotation marks ("").
- In the URL, each occurrence of the ampersand (&) character must be escaped with a backslash (\)

#### Example 27.3. Start a JBoss Data Grid Server with a Pre-Signed URL

```
bin/clustered.sh -Djboss.bind.address={server_ip_address} -
Djboss.bind.address.management={server_ip_address} -
Djboss.default.jgroups.stack=s3 -
Djgroups.s3.pre_signed_put_url="http://{s3_bucket_name}.s3.amazonaws.com
/ node1?
AWSAccessKeyId={access_key}&Expires={expiration_time}&Signature={signature}" -
Djgroups.s3.pre_signed_delete_url="http://{s3_bucket_name}.s3.amazonaws.
com/ node1?
AWSAccessKeyId={access_key}&Expires={expiration_time}&Signature={signature}"
```

In the provided example, the *{signatures}* values are generated by the **S3\_PING.generatePreSignedUrl()** method. Additionally, the *{expiration\_time}* values are the expiration time for the URL that are passed into the **S3\_PING.generatePreSignedUrl()** method.

[Report a bug](#)

### 27.2.3. Using Public S3 Buckets

This configuration involves an S3 bucket that has public read and write permissions, which means that **Everyone** has permissions to **List, Upload/Delete, View Permissions, and Edit Permissions** for the bucket.

The *location* property must be specified with the bucket name for this configuration. This configuration method is the least secure because any user who knows the name of the bucket can upload and store data in the bucket and the bucket creator's account is charged for this data.

To start the Red Hat JBoss Data Grid server, use the following command:

```
bin/clustered.sh -Djboss.bind.address={server_ip_address} -
Djboss.bind.address.management={server_ip_address} -
Djboss.default.jgroups.stack=s3 -Djgroups.s3.bucket={s3_bucket_name}
```

[Report a bug](#)

### 27.2.4. Troubleshooting S3\_PING Warnings

Depending on the **S3\_PING** configuration type used, the following warnings may appear when starting the JBoss Data Grid Server:

```
15:46:03,468 WARN [org.jgroups.conf.ProtocolConfiguration] (MSC service
thread 1-7) variable "${jgroups.s3.pre_signed_put_url}" in S3_PING could
not be substituted; pre_signed_put_url is removed from properties
```

```
15:46:03,469 WARN [org.jgroups.conf.ProtocolConfiguration] (MSC service
thread 1-7) variable "${jgroups.s3.prefix}" in S3_PING could not be
substituted; prefix is removed from properties
```

```
15:46:03,469 WARN [org.jgroups.conf.ProtocolConfiguration] (MSC service
thread 1-7) variable "${jgroups.s3.pre_signed_delete_url}" in S3_PING
could not be substituted; pre_signed_delete_url is removed from properties
```

■

In each case, ensure that the property listed as missing in the warning is not needed by the **S3\_PING** configuration.

[Report a bug](#)

# CHAPTER 28. USE RED HAT JBOSS DATA GRID WITH GOOGLE COMPUTE ENGINE

## 28.1. THE GOOGLE\_PING PROTOCOL

**GOOGLE\_PING** is a discovery protocol used by JGroups during cluster formation. It is ideal to use with Google Compute Engine (GCE) and uses Google Cloud Storage to store information about individual cluster members.

[Report a bug](#)

## 28.2. GOOGLE\_PING CONFIGURATION

Red Hat JBoss Data Grid works with Google Compute Engine in the following way:

- In Library mode, use the JGroups' configuration file **default-configs/default-jgroups-google.xml** or use the **GOOGLE\_PING** protocol in an existing configuration file.
- In Remote Client-Server mode, define the properties on the command line when you start the server to use the JGroups Google stack ( see example in [Section 28.2.1, “Starting the Server in Google Compute Engine”](#)).

To configure the **GOOGLE\_PING** protocol to work in Google Compute Engine in Library and Remote Client-Server mode:

- Use JGroups bucket. These buckets use Google Compute Engine credentials.
- Use the access key.
- Use the secret access key.



### NOTE

Only the **TCP** protocol is supported in Google Compute Engine since multicasts are not allowed.

[Report a bug](#)

### 28.2.1. Starting the Server in Google Compute Engine

This configuration requires access to a bucket that can only be accessed with the appropriate Google Compute Engine credentials.

Ensure that the **GOOGLE\_PING** configuration includes the following properties:

- the **access\_key** and the **secret\_access\_key** properties for the Google Compute Engine user.

#### Example 28.1. Start the Red Hat JBoss Data Grid Server with a Bucket

Run the following command from the top level of the server directory to start the Red Hat JBoss Data Grid server using a bucket:

■

```
bin/clustered.sh -Djboss.bind.address={server_ip_address} -  
Djboss.bind.address.management={server_ip_address} -  
Djboss.default.jgroups.stack=google -  
Djgroups.google.bucket={google_bucket_name} -  
Djgroups.google.access_key={access_key} -  
Djgroups.google.secret_access_key={secret_access_key}
```

1. Replace *{server\_ip\_address}* with the server's IP address.
2. Replace *{google\_bucket\_name}* with the appropriate bucket name.
3. Replace *{access\_key}* with the user's access key.
4. Replace *{secret\_access\_key}* with the user's secret access key.

[Report a bug](#)



## CHAPTER 29. HIGH AVAILABILITY USING SERVER HINTING

In Red Hat JBoss Data Grid, Server Hinting ensures that backed up copies of data are not stored on the same physical server, rack, or data center as the original. Server Hinting does not apply to total replication because total replication mandates complete replicas on every server, rack, and data center.

Data distribution across nodes is controlled by the Consistent Hashing mechanism. JBoss Data Grid offers a pluggable policy to specify the consistent hashing algorithm. For details see [Section 29.4, “ConsistentHashFactories”](#)

Setting a *machineId*, *rackId*, or *siteId* in the transport configuration will trigger the use of **TopologyAwareConsistentHashFactory**, which is the equivalent of the **DefaultConsistentHashFactory** with Server Hinting enabled.

Server Hinting is particularly important when ensuring the high availability of your JBoss Data Grid implementation.

[Report a bug](#)

### 29.1. ESTABLISHING SERVER HINTING WITH JGROUPS

When setting up a clustered environment in Red Hat JBoss Data Grid, Server Hinting is configured when establishing JGroups configuration.

JBoss Data Grid ships with several JGroups files pre-configured for clustered mode. These files can be used as a starting point when configuring Server Hinting in JBoss Data Grid.

**See Also:**

- [Section 26.2.2, “Pre-Configured JGroups Files”](#)

[Report a bug](#)

### 29.2. CONFIGURE SERVER HINTING (REMOTE CLIENT-SERVER MODE)

In Red Hat JBoss Data Grid's Remote Client-Server mode, Server Hinting is configured in the JGroups subsystem on the **transport** element for the default stack, as follows:

#### Procedure 29.1. Configure Server Hinting in Remote Client-Server Mode

```
<subsystem xmlns="urn:jboss:domain:jgroups:1.1"
  default-stack="{jboss.default.jgroups.stack:udp}">
  <stack name="udp">
    <transport type="UDP"
      socket-binding="jgroups-udp"
      site="{jboss.jgroups.transport.site:s1}"
      rack="{jboss.jgroups.transport.rack:r1}"
      machine="{jboss.jgroups.transport.machine:m1}">
      <!-- Additional configuration elements here -->
    </transport>
  </stack>
</subsystem>
```

1. Find the JGroups subsystem configuration
2. Enable Server Hinting via the **transport** Element
  - a. Set the site ID using the **site** parameter.
  - b. Set the rack ID using the **rack** parameter.
  - c. Set the machine ID using the **machine** parameter.

[Report a bug](#)

## 29.3. CONFIGURE SERVER HINTING (LIBRARY MODE)

In Red Hat JBoss Data Grid's Library mode, Server Hinting is configured at the transport level. The following is a Server Hinting sample configuration:

### Procedure 29.2. Configure Server Hinting for Library Mode

The following configuration attributes are used to configure Server Hinting in JBoss Data Grid.

```
<transport clusterName = "MyCluster"
           machineId = "LinuxServer01"
           rackId = "Rack01"
           siteId = "US-WestCoast" />
```

1. The **clusterName** attribute specifies the name assigned to the cluster.
2. The **machineId** attribute specifies the JVM instance that contains the original data. This is particularly useful for nodes with multiple JVMs and physical hosts with multiple virtual hosts.
3. The **rackId** parameter specifies the rack that contains the original data, so that other racks are used for backups.
4. The **siteId** parameter differentiates between nodes in different data centers replicating to each other.

The listed parameters are optional in a JBoss Data Grid configuration.

If **machineId**, **rackId**, or **siteId** are included in the configuration, **TopologyAwareConsistentHashFactory** is selected automatically, enabling Server Hinting. However, if Server Hinting is not configured, JBoss Data Grid's distribution algorithms are allowed to store replications in the same physical machine/rack/data center as the original data.

[Report a bug](#)

## 29.4. CONSISTENTHASHFACTORIES

Red Hat JBoss Data Grid offers a pluggable mechanism for selecting the consistent hashing algorithm. It is shipped with four implementations but a custom implementation can also be used.

JBoss Data Grid ships with four ConsistentHashFactory implementations:

- **DefaultConsistentHashFactory** - keeps segments balanced evenly across all the nodes, however the key mapping is not guaranteed to be same across caches, as this depends on the history of each cache. If no consistentHashFactory is specified this is the class that will be used.
- **SyncConsistentHashFactory** - guarantees that the key mapping is the same for each cache, provided the current membership is the same. This has a drawback in that a node joining the cache can cause the existing nodes to also exchange segments, resulting in either additional state transfer traffic, the distribution of the data becoming less even, or both.
- **TopologyAwareConsistentHashFactory** - equivalent of **DefaultConsistentHashFactory**, but automatically selected when the configuration includes server hinting.
- **TopologyAwareSyncConsistentHashFactory** - equivalent of **SyncConsistentHashFactory**, but automatically selected when the configuration includes server hinting.

The consistent hash implementation can be selected via the hash configuration:

```
<hash
consistentHashFactory="org.infinispan.distribution.ch.SyncConsistentHashFactory"/>
```

This configuration guarantees caches with the same members have the same consistent hash, and if the *machineId*, *rackId*, or *siteId* attributes are specified in the transport configuration it also spreads backup copies across physical machines/racks/data centers.

It has a potential drawback in that it can move a greater number of segments than necessary during re-balancing. This can be mitigated by using a larger number of segments.

Another potential drawback is that the segments are not distributed as evenly as possible, and actually using a very large number of segments can make the distribution of segments worse.

Despite the above potential drawbacks the **SyncConsistentHashFactory** and **TopologyAwareSyncConsistentHashFactory** both tend to reduce overhead in clustered environments, as neither of these calculate the hash based on the order that nodes have joined the cluster. In addition, both of these classes are typically faster than the default algorithms as both of these classes allow larger differences in the number of segments allocated to each node.

[Report a bug](#)

### 29.4.1. Implementing a ConsistentHashFactory

A custom **ConsistentHashFactory** must implement the **org.infinispan.distribution.ch.ConsistentHashFactory** interface with the following methods (all of which return an implementation of **org.infinispan.distribution.ch.ConsistentHash**):

#### Example 29.1. ConsistentHashFactory Methods

```
create(Hash hashFunction, int numOwners, int numSegments, List<Address>
members, Map<Address, Float> capacityFactors)
updateMembers(ConsistentHash baseCH, List<Address> newMembers,
Map<Address,
```

```
Float> capacityFactors)
rebalance(ConsistentHash baseCH)
union(ConsistentHash ch1, ConsistentHash ch2)
```

Currently it is not possible to pass custom parameters to **ConsistentHashFactory** implementations.

[Report a bug](#)

## 29.5. KEY AFFINITY SERVICE

The key affinity service allows a value to be placed in a certain node in a distributed Red Hat JBoss Data Grid cluster. The service returns a key that is hashed to a particular node based on a supplied cluster address identifying it.

The keys returned by the key affinity service cannot hold any meaning, such as a username. These are only random identifiers that are used throughout the application for this record. The provided key generators do not guarantee that the keys returned by this service are unique. For custom key format, you can pass your own implementation of **KeyGenerator**.

The following is an example of how to obtain and use a reference to this service.

### Example 29.2. Key Affinity Service

```
EmbeddedCacheManager cacheManager = getCacheManager();
Cache cache = cacheManager.getCache();
KeyAffinityService keyAffinityService =
    KeyAffinityServiceFactory.newLocalKeyAffinityService(
        cache,
        new RndKeyGenerator(),
        Executors.newSingleThreadExecutor(),
        100);
Object localKey =
    keyAffinityService.getKeyForAddress(cacheManager.getAddress());
cache.put(localKey, "yourValue");
```

The following procedure is an explanation of the provided example.

### Procedure 29.3. Using the Key Affinity Service

1. Obtain a reference to a cache manager and cache.
2. This starts the service, then uses the supplied **Executor** to generate and queue keys.
3. Obtain a key from the service which will be mapped to the local node (**cacheManager.getAddress()** returns the local address).
4. The entry with a key obtained from the **KeyAffinityService** is always stored on the node with the provided address. In this case, it is the local node.

[Report a bug](#)

### 29.5.1. Lifecycle

**KeyAffinityService** extends *Lifecycle*, which allows the key affinity service to be stopped, started, and restarted.

### Example 29.3. Key Affinity Service Lifecycle Parameter

```
public interface Lifecycle {
    void start();
    void stop();
}
```

The service is instantiated through the **KeyAffinityServiceFactory**. All factory methods have an **Executor**, that is used for asynchronous key generation, so that this does not occur in the caller's thread. The user controls the shutting down of this **Executor**.

The **KeyAffinityService** must be explicitly stopped when it is no longer required. This stops the background key generation, and releases other held resources. The **KeyAffinityService** will only stop itself when the cache manager with which it is registered is shut down.

[Report a bug](#)

## 29.5.2. Topology Changes

**KeyAffinityService** key ownership may change when a topology change occurs. The key affinity service monitors topology changes and updates so that it doesn't return stale keys, or keys that would map to a different node than the one specified. However, this does not guarantee that a node affinity hasn't changed when a key is used. For example:

1. Thread (**T1**) reads a key (**K1**) that maps to a node (**A**).
2. A topology change occurs, resulting in **K1** mapping to node **B**.
3. **T1** uses **K1** to add something to the cache. At this point, **K1** maps to **B**, a different node to the one requested at the time of read.

The above scenario is a not ideal, however it is a supported behavior for the application, as the keys that are already in use may be moved over during cluster change. The **KeyAffinityService** provides an access proximity optimization for stable clusters, which does not apply during the instability of topology changes.

[Report a bug](#)

## CHAPTER 30. SET UP CROSS-DATACENTER REPLICATION

In Red Hat JBoss Data Grid, Cross-Datcenter Replication allows the administrator to create data backups in multiple clusters. These clusters can be at the same physical location or different ones. JBoss Data Grid's Cross-Site Replication implementation is based on JGroups' **RELAY2** protocol.

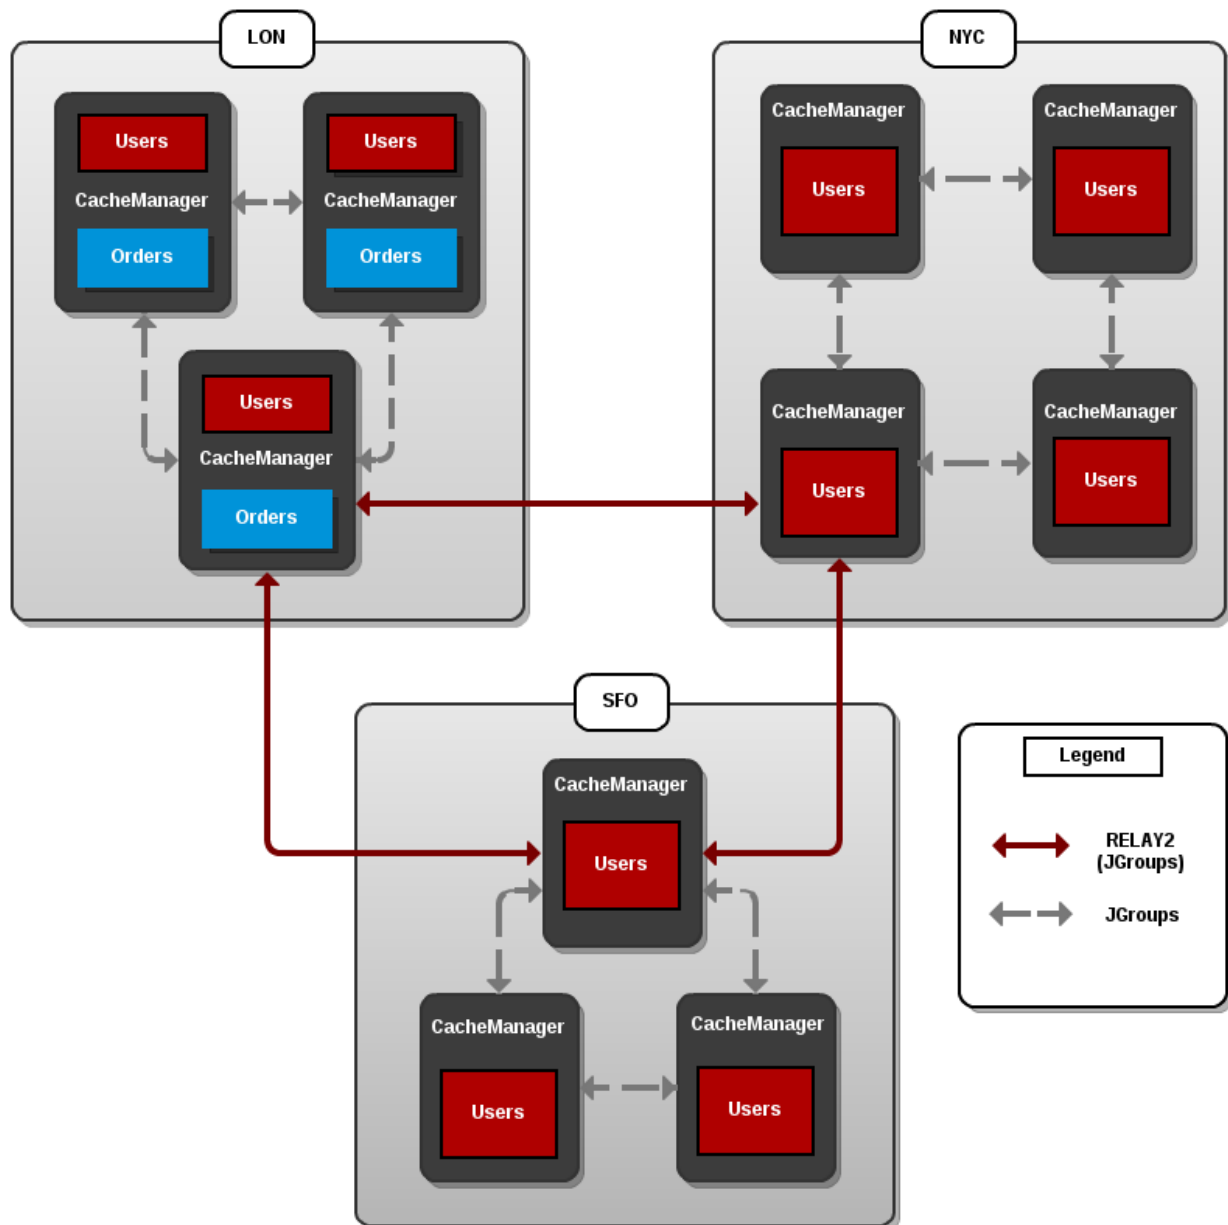
Cross-Datcenter Replication ensures data redundancy across clusters. Ideally, each of these clusters must be in a different physical location than the others.

[Report a bug](#)

### 30.1. CROSS-DATACENTER REPLICATION OPERATIONS

Red Hat JBoss Data Grid's Cross-Datcenter Replication operation is explained through the use of an example, as follows:

**Example 30.1. Cross-Datcenter Replication Example**



**Figure 30.1. Cross-Datcenter Replication Example**

Three sites are configured in this example: **LON**, **NYC** and **SFO**. Each site hosts a running JBoss Data Grid cluster made up of three to four physical nodes.

The **Users** cache is active in all three sites - **LON**, **NYC** and **SFO**. Changes to the **Users** cache at the any one of these sites will be replicated to the other two as long as the cache defines the other two sites as its backups through configuration. The **Orders** cache, however, is only available locally at the **LON** site because it is not replicated to the other sites.

The **Users** cache can use different replication mechanisms each site. For example, it can back up data synchronously to **SFO** and asynchronously to **NYC** and **LON**.

The **Users** cache can also have a different configuration from one site to another. For example, it can be configured as a distributed cache with *numOwners* set to **2** in the **LON** site, as a replicated cache in the **NYC** site and as a distributed cache with *numOwners* set to **1** in the **SFO** site.

JGroups is used for communication within each site as well as inter-site communication. Specifically, a JGroups protocol called **RELAY2** facilitates communication between sites. For more information, see [Section F.4, “About RELAY2”](#)

[Report a bug](#)

## 30.2. CONFIGURE CROSS-DATACENTER REPLICATION

### 30.2.1. Configure Cross-Datcenter Replication (Remote Client-Server Mode)

In Red Hat JBoss Data Grid's Remote Client-Server mode, cross-datcenter replication is set up as follows:

#### Procedure 30.1. Set Up Cross-Datcenter Replication

##### 1. Set Up RELAY

Add the following configuration to the `standalone.xml` file to set up **RELAY**:

```
<subsystem xmlns="urn:jboss:domain:jgroups:1.2"
  default-stack="udp">
  <stack name="udp">
    <transport type="UDP"
      socket-binding="jgroups-udp"/>
    <!-- Additional configuration elements here -->
    <relay site="LON">
      <remote-site name="NYC" stack="tcp" cluster="global"/>
      <remote-site name="SFO" stack="tcp" cluster="global"/>
      <property name="relay_multicasts">false</property>
    </relay>
  </stack>
</subsystem>
```

The **RELAY** protocol creates an additional stack (running parallel to the existing **TCP** stack) to communicate with the remote site. If a **TCP** based stack is used for the local cluster, two **TCP** based stack configurations are required: one for local communication and one to connect to the remote site. For an illustration, see [Section 30.1, “Cross-Datcenter Replication Operations”](#)

## 2. Set Up Sites

Use the following configuration in the `standalone.xml` file to set up sites for each distributed cache in the cluster:

```
<distributed-cache>
  <!-- Additional configuration elements here -->
  <backups>
    <backup site="{FIRSTSITEName}" strategy="{SYNC/ASync}" />
    <backup site="{SECONDSITENAME}" strategy="{SYNC/ASync}" />
  </backups>
</distributed-cache>
```

## 3. Configure Local Site Transport

Add the name of the local site in the `transport` element to configure transport:

```
<transport executor="infinispan-transport"
  lock-timeout="60000"
  cluster="LON"
  stack="udp"/>
```

[Report a bug](#)

## 30.2.2. Configure Cross-Data Replication (Library Mode)

### 30.2.2.1. Configure Cross-Datacenter Replication Declaratively

When configuring Cross-Datacenter Replication, the `relay.RELAY2` protocol creates an additional stack (running parallel to the existing `TCP` stack) to communicate with the remote site. If a `TCP`-based stack is used for the local cluster, two `TCP` based stack configurations are required: one for local communication and one to connect to the remote site.

In JBoss Data Grid's Library mode, cross-datacenter replication is set up as follows:

### Procedure 30.2. Setting Up Cross-Datacenter Replication

#### 1. Configure the Local Site

```
<infinispan>
  <global>
    <site local="SF0" />
    <transport clusterName="default">
      <properties>
        <property name="configurationFile" value="jgroups-
with-relay.xml"/>
      </properties>
    </transport>
    <!-- Additional configuration information here -->
  </global>
  <!-- Additional configuration information here -->
  <namedCache name="lonBackup">
    <sites>
      <backupFor remoteSite="LON"
remoteCache="lon" />
    </sites>
```



```

    </namedCache>
</infinispan>

```

- a. Add the **site** element to the **global** element to add the local site (in this example, the local site is named **LON**).
- b. Cross-site replication requires a non-default JGroups configuration. Add the **transport** element and set up the path to the configuration file as the **configurationFile** property. In this example, the JGroups configuration file is named **jgroups-with-relay.xml**.
- c. Configure the cache in site **LON** to back up to the sites **NYC** and **SFO**.
- d. Configure the back up caches:
  - i. Configure the cache in site **NYC** to receive back up data from **LON**.
  - ii. Configure the cache in site **SFO** to receive back up data from **LON**.

## 2. Add the Contents of the Configuration File

As a default, Red Hat JBoss Data Grid includes JGroups configuration files such as **default-configs/default-jgroups-tcp.xml** and **default-configs/default-jgroups-udp.xml** in the **infinispan-embedded-*{VERSION}*.jar** package.

Copy the JGroups configuration to a new file (in this example, it is named **jgroups-with-relay.xml**) and add the provided configuration information to this file. Note that the **relay.RELAY2** protocol configuration must be the last protocol in the configuration stack.

```

<config>
  ...
  <relay.RELAY2 site="LON"
    config="relay.xml"
    relay_multicasts="false" />
</config>

```

## 3. Configure the relay.xml File

Set up the **relay.RELAY2** configuration in the **relay.xml** file. This file describes the global cluster configuration.

```

<RelayConfiguration>
  <sites>
    <site name="LON"
      id="0">
      <bridges>
        <bridge config="jgroups-global.xml"
          name="global"/>
      </bridges>
    </site>
    <site name="NYC"
      id="1">
      <bridges>
        <bridge config="jgroups-global.xml"
          name="global"/>
      </bridges>
    </site>
    <site name="SFO"

```

```

        id="2">
      <bridges>
        <bridge config="jgroups-global.xml"
              name="global"/>
      </bridges>
    </site>
  </sites>
</RelayConfiguration>

```

#### 4. Configure the Global Cluster

The file `jgroups-global.xml` referenced in `relay.xml` contains another JGroups configuration which is used for the global cluster: communication between sites.

The global cluster configuration is usually **TCP**-based and uses the **TCPPING** protocol (instead of **PING** or **MPING**) to discover members. Copy the contents of `default-configs/default-jgroups-tcp.xml` into `jgroups-global.xml` and add the following configuration in order to configure **TCPPING**:

```

<config>
  <TCP bind_port="7800" ... />
  <TCPPING
    initial_hosts="lon.hostname[7800],nyc.hostname[7800],sfo.hostname[78
00]"
    ergonomics="false" />
  <!-- Rest of the protocols -->
</config>

```

Replace the hostnames (or IP addresses) in `TCPPING.initial_hosts` with those used for your site masters. The ports (**7800** in this example) must match the `TCP.bind_port`.

For more information about the **TCPPING** protocol, see [Section 26.2.1.3, "Using the TCPping Protocol"](#)

[Report a bug](#)

### 30.2.2.2. Configure Cross-Datacenter Replication Programmatically

The programmatic method to configure cross-datacenter replication in Red Hat JBoss Data Grid is as follows:

#### Procedure 30.3. Configure Cross-Datacenter Replication Programmatically

##### 1. Identify the Node Location

Declare the site the node resides in:

```
globalConfiguration.site().localSite("LON");
```

##### 2. Configure JGroups

Configure JGroups to use the **RELAY** protocol:

```
globalConfiguration.transport().addProperty("configurationFile",
jgroups-with-relay.xml);
```


### 3. Set Up the Remote Site

Set up JBoss Data Grid caches to replicate to the remote site:

```
ConfigurationBuilder lon = new ConfigurationBuilder();
lon.sites().addBackup()
    .site("NYC")
    .backupFailurePolicy(BackupFailurePolicy.WARN)
    .strategy(BackupConfiguration.BackupStrategy.SYNC)
    .replicationTimeout(12000)
    .sites().addInUseBackupSite("NYC")
.sites().addBackup()
    .site("SFO")
    .backupFailurePolicy(BackupFailurePolicy.IGNORE)
    .strategy(BackupConfiguration.BackupStrategy.ASYNC)
    .sites().addInUseBackupSite("SFO")
```

### 4. Optional: Configure the Backup Caches

JBoss Data Grid implicitly replicates data to a cache with same name as the remote site. If a backup cache on the remote site has a different name, users must specify a **backupFor** cache to ensure data is replicated to the correct cache.



#### NOTE

This step is optional and only required if the remote site's caches are named differently from the original caches.

- a. Configure the cache in site **NYC** to receive backup data from **LON**:

```
ConfigurationBuilder NYCbackupOfLon = new ConfigurationBuilder();
NYCbackupOfLon.sites().backupFor().remoteCache("lon").remoteSite(
    "LON");
```

- b. Configure the cache in site **SFO** to receive backup data from **LON**:

```
ConfigurationBuilder SFObackupOfLon = new ConfigurationBuilder();
SFObackupOfLon.sites().backupFor().remoteCache("lon").remoteSite(
    "LON");
```

### 5. Add the Contents of the Configuration File

As a default, Red Hat JBoss Data Grid includes JGroups configuration files such as **default-configs/default-jgroups-tcp.xml** and **default-configs/default-jgroups-udp.xml** in the **infinispan-embedded-*{VERSION}*.jar** package.

Copy the JGroups configuration to a new file (in this example, it is named **jgroups-with-relay.xml**) and add the provided configuration information to this file. Note that the **relay.RELAY2** protocol configuration must be the last protocol in the configuration stack.

```
<config>
  <!-- Additional configuration information here -->
  <relay.RELAY2 site="LON"
    config="relay.xml"
    relay_multicasts="false" />
</config>
```

## 6. Configure the relay.xml File

Set up the **relay.RELAY2** configuration in the **relay.xml** file. This file describes the global cluster configuration.

```
<RelayConfiguration>
  <sites>
    <site name="LON"
      id="0">
      <bridges>
        <bridge config="jgroups-global.xml"
          name="global"/>
      </bridges>
    </site>
    <site name="NYC"
      id="1">
      <bridges>
        <bridge config="jgroups-global.xml"
          name="global"/>
      </bridges>
    </site>
    <site name="SFO"
      id="2">
      <bridges>
        <bridge config="jgroups-global.xml"
          name="global"/>
      </bridges>
    </site>
  </sites>
</RelayConfiguration>
```

## 7. Configure the Global Cluster

The file **jgroups-global.xml** referenced in **relay.xml** contains another JGroups configuration which is used for the global cluster: communication between sites.

The global cluster configuration is usually **TCP**-based and uses the **TCPPING** protocol (instead of **PING** or **MPING**) to discover members. Copy the contents of **default-configs/default-jgroups-tcp.xml** into **jgroups-global.xml** and add the following configuration in order to configure **TCPPING**:

```
<config>
  <TCP bind_port="7800" <!-- Additional configuration information
  here --> />
  <TCPPING
  initial_hosts="lon.hostname[7800],nyc.hostname[7800],sfo.hostname[78
  00]"
    ergonomics="false" />
  <!-- Rest of the protocols -->
</config>
```

Replace the hostnames (or IP addresses) in **TCPPING.initial\_hosts** with those used for your site masters. The ports (**7800** in this example) must match the **TCP.bind\_port**.

For more information about the **TCPPING** protocol, see [Section 26.2.1.3, “Using the TCPping Protocol”](#)

[Report a bug](#)

## 30.3. TAKING A SITE OFFLINE

In Red Hat JBoss Data Grid's Cross-datacenter replication configuration, if backing up to one site fails a certain number of times during a time interval, that site can be marked as offline automatically. This feature removes the need for manual intervention by an administrator to mark the site as offline.

It is possible to configure JBoss Data Grid to take down a site automatically when specified conditions are met, or for an administrator to manually take down a site:

- Configure automatically taking a site offline:
  - Declaratively in Remote Client-Server mode.
  - Declaratively in Library mode.
  - Using the programmatic method.
- Manually taking a site offline:
  - Using JBoss Operations Network (JON).
  - Using the JBoss Data Grid Command Line Interface (CLI).

[Report a bug](#)

### 30.3.1. Taking a Site Offline (Remote Client-Server Mode)

In Red Hat JBoss Data Grid's Remote Client-Server mode, the **take-offline** element is added to the **backup** element to configure when a site is automatically taken offline.

#### Example 30.2. Taking a Site Offline in Remote Client-Server Mode

```
<backup>
  <take-offline after-failures="{NUMBER}"
    min-wait="{PERIOD}" />
</backup>
```

The **take-offline** element use the following parameters to configure when to take a site offline:

- The **after-failures** parameter specifies the number of times attempts to contact a site can fail before the site is taken offline.
- The **min-wait** parameter specifies the number (in milliseconds) to wait to mark an unresponsive site as offline. The site is offline when the **min-wait** period elapses after the first attempt, and the number of failed attempts specified in the **after-failures** parameter occur.

[Report a bug](#)

### 30.3.2. Taking a Site Offline (Library Mode)

In Red Hat JBoss Data Grid's Library mode, use the **backupFor** element after defining all back up sites within the **backups** element:

#### Example 30.3. Taking a Site Offline in Library Mode

```
<backup>
  <takeOffline afterFailures="{NUM}"
              minTimeToWait="{PERIOD}"/>
</backup>
```

Add the **takeOffline** element to the **backup** element to configure automatically taking a site offline.

- The **afterFailures** parameter specifies the number of times attempts to contact a site can fail before the site is taken offline. The default value (0) allows an infinite number of failures if **minTimeToWait** is less than 0. If the **minTimeToWait** is not less than 0, **afterFailures** behaves as if the value is negative. A negative value for this parameter indicates that the site is taken offline after the time specified by **minTimeToWait** elapses.
- The **minTimeToWait** parameter specifies the number (in milliseconds) to wait to mark an unresponsive site as offline. The site is taken offline after the number attempts specified in the **afterFailures** parameter conclude and the time specified by **minTimeToWait** after the first failure has elapsed. If this parameter is set to a value smaller than or equal to 0, this parameter is disregarded and the site is taken offline based solely on the **afterFailures** parameter.

[Report a bug](#)

### 30.3.3. Taking a Site Offline (Programmatically)

To configure taking a Cross-datacenter replication site offline automatically in Red Hat JBoss Data Grid programmatically:

#### Example 30.4. Taking a Site Offline Programmatically

```
lon.sites().addBackup()
    .site("NYC")
    .backupFailurePolicy(BackupFailurePolicy.FAIL)
    .strategy(BackupConfiguration.BackupStrategy.SYNC)
    .takeOffline()
        .afterFailures(500)
        .minTimeToWait(10000);
```

[Report a bug](#)

### 30.3.4. Taking a Site Offline via JBoss Operations Network (JON)

A site can be taken offline in Red Hat JBoss Data Grid using the JBoss Operations Network operations. For a list of the metrics, see [Section 22.6.2, “JBoss Operations Network Plugin Operations”](#)

[Report a bug](#)

### 30.3.5. Taking a Site Offline via the CLI

Use Red Hat JBoss Data Grid's Command Line Interface (CLI) to manually take a site from a cross-datacenter replication configuration down if it is unresponsive using the **site** command.

The **site** command can be used to check the status of a site as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> site --status
${SITENAME}
```

The result of this command would either be **online** or **offline** according to the current status of the named site.

The command can be used to bring a site online or offline by name as follows:

```
[jmx://localhost:12000/MyCacheManager/namedCache]> site --offline
${SITENAME}
```

```
[jmx://localhost:12000/MyCacheManager/namedCache]> site --online
${SITENAME}
```

If the command is successful, the output **ok** displays after the command. As an alternate, the site can also be brought online using JMX (see [Section 30.3.6, “Bring a Site Back Online”](#) for details).

For more information about the JBoss Data Grid CLI and its commands, see the *Developer Guide's* chapter on the JBoss Data Grid Command Line Interface (CLI)

[Report a bug](#)

### 30.3.6. Bring a Site Back Online

After a site is taken offline, the site can be brought back online either using the JMX console to invoke the **bringSiteOnline(siteName)** operation on the **XSiteAdmin** MBean (See [Section C.23, “XSiteAdmin”](#) for details) or using the CLI (see [Section 30.3.5, “Taking a Site Offline via the CLI”](#) for details).

[Report a bug](#)

## 30.4. STATE TRANSFER BETWEEN SITES

When an offline master site is back online, it is necessary to synchronize its state with the latest data from the backup site. State transfer allows state to be transferred from one site to another, meaning the master site is synchronized and made consistent with the backup site. Similarly, when a backup site becomes available, state transfer can be utilized to make it consistent with the master site.

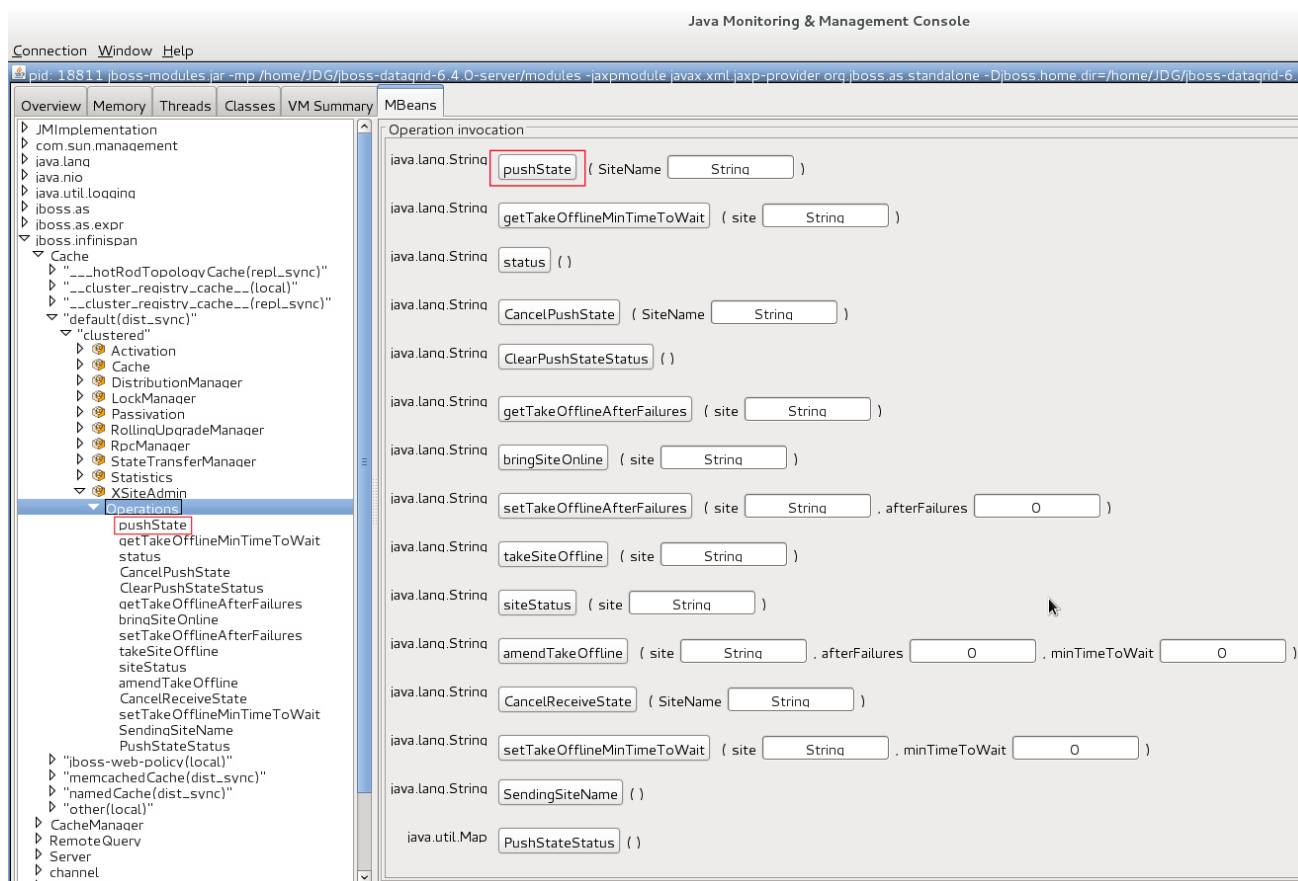
Consider a scenario of two sites - Master site A and Backup site B. Clients can originally access only Master site A whereas Backup Site B acts as an invisible backup. Cross Site State Transfer can be pushed bidirectionally. When the new backup site B goes online, in order to synchronize its state with the master site A, a State Transfer can be initiated to push the state from the Master site A to the Backup site B.

Similarly, when the Master site A is brought back online, in order to synchronize it with the Backup site B, a State Transfer can be initiated to push the state from Backup site B to Master Site A.

The use cases applies for both Active-Passive and Active-Active State Transfer. The difference is that during Active-Active State Transfer we assume that cache operations can be performed in the site, which consumes state.

A system administrator or an authorized entity initiates the state transfer manually using JMX. The system administrator invokes the **pushState(SiteName String)** operation available in the **XSiteAdminOperations** MBean.

The following interface shows the **pushState(SiteName String)** operation in JConsole:



**Figure 30.2. PushState Operation**

State transfer is also invoked using the Command Line Interface (CLI) by the **site push sitename** command. For example, when the master site is brought back online, the system administrator invokes the state transfer operation in the backup site, specifying the master site name that is to receive the state.

The master site can be offline at the time of the push operation. On successful state transfer, the state data common to both the sites is overwritten on the master site. For example, if key A exists on the master site but not on the backup site, key A will not be deleted from the master site. Whereas, if key B exists on the backup as well as the master site, key B is overwritten on the master site.



#### NOTE

Updates on keys performed after initiating state transfer are not overwritten by incoming state transfer.

Cross-site state transfer can be transactional and supports 1PC and 2PC transaction options. 1PC and 2PC options define whether data modified inside a transaction is backed up to a remote site in one or two phases. 2PC includes a prepare phase in which backup sites acknowledges that transaction has been



successfully prepared. Both options are supported.

[Report a bug](#)

### 30.4.1. Active-Passive State Transfer

The active-passive state transfer is used when cross-site replication is used to back up the master site. The master site processes all the requests but if it goes offline, the backup site starts to handle them. When the master site is back online, it receives the state from the backup site and starts to handle the client requests. In Active-Passive state transfer mode, transactional writes happen concurrently with state transfer on the site which sends the state.

In active-passive state transfer mode, the client read-write requests occurs only on the backup site. The master site acts as an invisible backup until the client requests are switched to it when the state transfer is completed. The active-passive state transfer mode is fully supported in cross-datacenter replication.

When an Active-Passive State Transfer is interrupted by a network failure, the System Administrator invokes the JMX operation manually to resume the state transfer. To transfer the state, for example from Master site A to Backup site B, invoke the JMX operation on Master site A. Similarly, to transfer state from Backup site B to Master site A, invoke the JMX operation on the Backup site B.

The JMX operation is invoked on the site from which the state is transferred to the other site that is online to synchronize the states.

For example, there is a running backup site and the system administrator wants to bring back the master site online. To use active-passive state transfer, the system administrator will perform the following steps.

- Boot the Red Hat JBoss Data Grid cluster in the master site.
- Command the backup site to push state to the master site.
- Wait until the state transfer is complete.
- Make the clients aware that the master site is available to process the requests.

[Report a bug](#)

### 30.4.2. Active-Active State Transfer

In active-active state transfer mode, the client requests occur concurrently in both the sites while the state transfer is in progress. The current implementation supports handling requests in the new site while the state transfer is in progress, which may break the data consistency.



#### WARNING

Active-active state transfer mode is not fully supported, as it may lead to data inconsistencies.

**NOTE**

In active-active state transfer mode, both the sites, the master and the backup sites share the same role. There is no clear distinction between the master and backup sites in the active-active state transfer mode

For example, there is a running site and the system administrator wants to bring a new site online. To use active-active state transfer, the system administrator must perform the following steps.

- Boot the Red Hat JBoss Data Grid cluster in the new site.
- Command the running site to push state to the new site.
- Make the clients aware that the new site is available to process the requests.

[Report a bug](#)

### 30.4.3. State Transfer Configuration

State transfer between sites is not enabled or disabled but it allows to tune some parameters. The only configuration is done by the system administrator while configuring the load balancer to switch the request to the master site during or after the state transfer. The implementation handles a case in which a key is updated by a client before it receives the state, ignoring when it is delivered.

The following are default parameter values:

```
<backups>
  <backup site="NYC"
    strategy="SYNC"
    failure-policy="FAIL">
    <state-transfer chunk-size="512"
      timeout="1200000"
      max-retries="30"
      wait-time="2000" />
  </backup>
</backups>
```

[Report a bug](#)

## 30.5. CONFIGURE MULTIPLE SITE MASTERS

A standard Red Hat JBoss Data Grid cross-datacenter replication configuration includes one master node for each site. The master node is a gateway for other nodes to communicate with the master nodes at other sites.

This standard configuration works for a simple cross-datacenter replication configuration. However, with a larger volume of traffic between the sites, passing traffic through a single master node can create a bottleneck, which slows communication across nodes.

In JBoss Data Grid, configure multiple master nodes for each site to optimize traffic across multiple sites.

[Report a bug](#)

### 30.5.1. Multiple Site Master Operations

When multiple site masters are enabled and configured, the master nodes in each site joins the local cluster (i.e. the local site) as well as the global cluster (which includes nodes that are members of multiple sites).

Each node that acts as a site master and maintains a routing table that consists of a list of target sites and site masters. When a message arrives, a random master node for the destination site is selected. The message is then forwarded to the random master node, where it is sent to the destination node (unless the randomly selected node was the destination).

[Report a bug](#)

### 30.5.2. Configure Multiple Site Masters (Remote Client-Server Mode)

#### Prerequisites

Configure Cross-Datacenter Replication for Red Hat JBoss Data Grid's Remote Client-Server Mode.

#### Procedure 30.4. Set Multiple Site Masters in Remote Client-Server Mode

```
<relay site="LON">
  <remote-site name="NYC" stack="tcp" cluster="global"/>
  <remote-site name="SFO" stack="tcp" cluster="global"/>
  <property name="relay_multicasts">false</property>
  <property name="max_site_masters">16</property>
  <property name="can_become_site_master">true</property>
</relay>
```

##### 1. Locate the Target Configuration

Locate the target site's configuration in the `clustered-xsite.xml` example configuration file. The sample configuration looks like example provided above.

##### 2. Configure Maximum Sites

Use the `max_site_masters` property to determine the maximum number of master nodes within the site. Set this value to the number of nodes in the site to make every node a master.

##### 3. Configure Site Master

Use the `can_become_site_master` property to allow the node to become the site master. This flag is set to `true` as a default. Setting this flag to `false` prevents the node from becoming a site master. This is required in situations where the node does not have a network interface connected to the external network.

[Report a bug](#)

### 30.5.3. Configure Multiple Site Masters (Library Mode)

To configure multiple site masters in Red Hat JBoss Data Grid's Library Mode:

#### Procedure 30.5. Configure Multiple Site Masters (Library Mode)

##### 1. Configure Cross-Datacenter Replication

Configure Cross-Datacenter Replication in JBoss Data Grid. Use the instructions in [Section 30.2.2.1, "Configure Cross-Datacenter Replication Declaratively"](#) for an XML configuration or the instructions in [Section 30.2.2.2, "Configure Cross-Datacenter Replication Programmatically"](#) for a programmatic configuration.

## 2. Add the Contents of the Configuration File

Add the *can\_become\_site\_master* and *max\_site\_masters* parameters to the configuration as follows:

```
<config>
  <!-- Additional configuration information here -->
  <relay.RELAY2 site="LON"
    config="relay.xml"
    relay_multicasts="false"
    can_become_site_master="true"
    max_site_masters="16"/>
</config>
```

Set the *max\_site\_masters* value to the number of nodes in the cluster to make all nodes masters.

[Report a bug](#)

## CHAPTER 31. EXTERNALIZE SESSIONS

### 31.1. EXTERNALIZE HTTP SESSION FROM JBOSS EAP 6.4 AND LATER TO JBOSS DATA GRID

Red Hat JBoss Data Grid can be used as an external cache container for application specific data in JBoss Enterprise Application Platform (EAP), such as HTTP Sessions. This allows scaling of the data layer independent of the application, and enables different EAP clusters, that may reside in various domains, to access data from the same JBoss Data Grid cluster. Additionally, other applications can interface with the caches presented by Red Hat JBoss Data Grid.



#### NOTE

The following procedures have been tested and confirmed to function on JBoss EAP 6.4 and JBoss Data Grid 6.5. Externalizing HTTP Sessions should only be used on these, or later, versions of each product.

The below procedure applies for both standalone and domain mode of EAP; however, in domain mode each server group requires a unique remote cache configured. While multiple server groups can utilize the same Red Hat JBoss Data Grid cluster the respective remote caches will be unique to the EAP server group.



#### NOTE

For each distributable application, an entirely new cache must be created. It can be created in an existing cache container, for example, web.

#### Procedure 31.1. Externalize HTTP Sessions

1. Ensure the remote cache containers are defined in EAP's **infinispan** subsystem; in the example below the **cache** attribute in the **remote-store** element defines the cache name on the remote JBoss Data Grid server:

```
<subsystem xmlns="urn:jboss:domain:infinispan:1.5">
  [...]
  <cache-container name="cacheContainer" default-cache="default-cache" module="org.jboss.as.clustering.web.infinispan" statistics-enabled="true">
    <transport lock-timeout="60000"/>
    <replicated-cache name="default-cache" mode="SYNC" batching="true">
      <remote-store cache="default" socket-timeout="60000" preload="true" passivation="false" purge="false" shared="true">
        <remote-server outbound-socket-binding="remote-jdg-server1"/>
        <remote-server outbound-socket-binding="remote-jdg-server2"/>
      </remote-store>
    </replicated-cache>
  </cache-container>
</subsystem>
```

- Define the location of the remote Red Hat JBoss Data Grid server by adding the networking information to the **socket-binding-group**:

```
<socket-binding-group ...>
  <outbound-socket-binding name="remote-jdg-server1">
    <remote-destination host="JDGHostName1" port="11222"/>
  </outbound-socket-binding>
  <outbound-socket-binding name="remote-jdg-server2">
    <remote-destination host="JDGHostName2" port="11222"/>
  </outbound-socket-binding>
</socket-binding-group>
```

- Repeat the above steps for each cache-container and each Red Hat JBoss Data Grid server. Each server defined must have a separate **<outbound-socket-binding>** element defined.
- Add passivation and cache information into the application's **jboss-web.xml**. In the following example **cacheContainer** is the name of the cache container, and **default-cache** is the name of the default cache located in this container. An example file is shown below:

```
<?xml version="1.0" encoding="UTF-8"?>

<jboss-web version="6.0"
  xmlns="http://www.jboss.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.jboss.com/xml/ns/javaee
http://www.jboss.org/j2ee/schema/jboss-web_6_0.xsd">

  <replication-config>
    <replication-trigger>SET</replication-trigger>
    <replication-granularity>SESSION</replication-granularity>
    <cache-name>cacheContainer.default-cache</cache-name>
  </replication-config>

  <passivation-config>
    <use-session-passivation>true</use-session-passivation>
    <passivation-min-idle-time>900</passivation-min-idle-time>
    <passivation-max-idle-time>1800</passivation-max-idle-time>
  </passivation-config>

</jboss-web>
```



#### NOTE

The passivation timeouts above are provided assuming that a typical session is abandoned within 15 minutes and uses the default HTTP session timeout in JBoss EAP of 30 minutes. These values may need to be adjusted based on each application's workload.

[Report a bug](#)

## CHAPTER 32. HANDLING NETWORK PARTITIONS (SPLIT BRAIN)

Network Partitions occur when a cluster breaks into two or more partitions. As a result, the nodes in each partition are unable to locate or communicate with nodes in the other partitions. This results in an unintentionally partitioned network.

In the event of a network partition in a distributed system like Red Hat JBoss Data Grid, the CAP (Brewer's) theorem comes into play. The CAP theorem states that in the event of a Network Partition (P), a distributed system can provide either Consistency (C) or Availability (A) for the data, but not both.

By default, Partition Handling is disabled in JBoss Data Grid. During a network partition, the partitions continue to remain Available (A), at the cost of Consistency (C).

However, when Partition Handling is enabled, JBoss Data Grid prioritizes consistency (C) of data over Availability (A).

Red Hat JBoss Data Grid offers the primary partitions strategy to repair a split network. When the network partition occurs and the cache is split into two or more partitions, at most one partition becomes the primary partition (and stays available) and the others are designated as secondary partitions (and enter Degraded Mode). When the partitions merge back into a single cache, the primary partition is then used as a reference for all secondary partitions. All members of the secondary partitions must remove their current state information and replace it with fresh state information from a member of the primary partition. If there was no primary partition during the split, the state on every node is assumed to be correct.

In JBoss Data Grid, a cache consists of data stored on a number of nodes. To prevent data loss if a node fails, JBoss Data Grid replicates a data item over multiple nodes. In distribution mode, this redundancy is configured using the *numOwners* configuration attribute, which specifies the number of replicas for each cache entry in the cache. As a result, as long as the number of nodes that have failed are less than the value of *numOwners*, JBoss Data Grid retains a copy of the lost data and can recover.



### NOTE

In JBoss Data Grid's replication mode, however, *numOwners* is always equal to the number of nodes in the cache, because each node contains a copy of every data item in the cache in this mode.

In certain cases, a number of nodes greater than the value of *numOwners* can disappear from the cache. Two common reasons for this are:

- **Split-Brain:** Usually, as the result of a router crash, the cache is divided into two or more partitions. Each of the partitions operates independently of the other and each may contain different versions of the same data.
- **Successive Crashed Nodes:** A number of nodes greater than the value of *numOwners* crashes in succession for any reason. JBoss Data Grid is unable to properly balance the state between crashes, and the result is partial data loss.

[Report a bug](#)

### 32.1. DETECTING AND RECOVERING FROM A SPLIT-BRAIN PROBLEM

When a Split-Brain occurs in the data grid, each network partition installs its own JGroups view with

nodes from other partitions removed. The partitions remain unaware of each other, therefore there is no way to determine how many partitions the network has split into. Red Hat JBoss Data Grid assumes that the cache has unexpectedly split if one or more nodes disappear from the JGroups cache without sending an explicit leaving message, while in reality the cause can be physical (crashed switches, cable failure, etc.) to virtual (stop-the-world garbage collection).

This state is dangerous because each of the newly split partitions operates independently and can store conflicting updates for the same data entries.

When Partition Handling mode is enabled (see [Section 32.6, “Configure Partition Handling”](#) for instructions) and JBoss Data Grid suspects that one or more nodes are no longer accessible, each partition does not start a rebalance immediately, but first it checks whether it should enter degraded mode instead. To enter Degraded Mode, one of the following conditions must be true:

- At least one segment has lost all its owners, which means that a number of nodes equal to or greater than the value of *numOwners* have left the JGroups view.
- The partition does not contain a majority of nodes (greater than half) of the nodes from the latest stable topology. The stable topology is updated each time a rebalance operation successfully concludes and the coordinator determines that additional rebalancing is not required.

If neither of the conditions are met, the partition continues normal operations and JBoss Data Grid attempts to rebalance its nodes. Based on these conditions, at most one partition can remain in Available mode. Other partitions will enter Degraded Mode.

When a partition enters into Degraded Mode, it only allows read/write access to those entries for which all owners (copies) of the entry exist on nodes within the same partition. Read and write requests for an entry for which one or more of its owners (copies) exist on nodes that have disappeared from the partition are rejected with an **AvailabilityException**.



#### NOTE

A possible limitation is that if two partitions start as isolated partitions and do not merge, they can read and write inconsistent data. JBoss Data Grid does not identify such partitions as split partitions.





## WARNING

Data consistency can be at risk from the time (t1) when the cache physically split to the time (t2) when JBoss Data Grid detects the connectivity change and changes the state of the partitions:

- Transactional writes that were in progress at t1 when the split physically occurred may be rolled back on some of the owners. This can result in inconsistency between the copies (after the partitions rejoin) of an entry that is affected by such a write. However, transactional writes that started after t1 will fail as expected.
- If the write is non-transactional, then during this time window, a value written only in a minor partition (due to physical split and because the partition has not yet been Degraded) can be lost when partitions rejoin, if this minor partition receives state from a primary (Available) partition upon rejoin. If the partition does not receive state upon rejoin (i.e. all partitions are degraded), then the value is not lost, but an inconsistency can remain.
- There is also a possibility of a stale read in a minor partition during this transition period, as an entry is still Available until the minor partition enters Degraded state.

When partitions merge after a network partition has occurred:

- If one of the partitions was Available during the network partition, then the joining partition(s) are wiped out and state transfer occurs from the Available (primary) partition to the joining nodes.
- If all joining partitions were Degraded during the Split Brain, then no state transfer occurs during the merge. The combined cache is then Available only if the merging partitions contain a simple majority of the members in the latest stable topology (one with the highest topology ID) and has at least an owner for each segment (i.e. keys are not lost).



## WARNING

Between the time (t1) when partitions begin merging to the time (t2) when the merge is complete, nodes reconnect through a series of merge events. During this time window, it is possible that a node can be reported as having temporarily left the cluster. For a Transactional cache, if during this window between t1 and t2, such a node is executing a transaction that spans other nodes, then this transaction may not execute on the remote node, but still succeed on the originating node. The result is a potential stale value for affected entries on a node that did not commit this transaction.


After t2, once the merge has completed on all nodes, this situation will not occur for subsequent transactions. However, an inconsistency introduced on entries that were affected by a transaction in progress during the time window between t1 and t2 is not resolved until these entries are subsequently updated or deleted. Until then, a read on such impacted entries can potentially return the stale value.

[Report a bug](#)

## 32.2. SPLIT BRAIN TIMING: DETECTING A SPLIT

When using the **FD\_ALL** protocol a given node becomes suspected after the following amount of milliseconds have passed:

```
FD_ALL.timeout + FD_ALL.interval + VERIFY_SUSPECT.timeout +
GMS.view_ack_collection_timeout
```



## IMPORTANT

The amount of time taken in the formulas above is how long it takes JBoss Data Grid to install a cluster view without the leavers; however, as JBoss Data Grid runs inside a JVM excessive Garbage Collection (GC) times can increase this time beyond the failure detection outlined above. JBoss Data Grid has no control over these GC times, and excessive GC on the coordinator can delay this detection by an amount equal to the GC time.

[Report a bug](#)

## 32.3. SPLIT BRAIN TIMING: RECOVERING FROM A SPLIT

After a split occurs JBoss Data Grid will merge the partitions back, and the maximum time to detect a merge after the network partition is healed is:

```
3.1 * MERGE3.max_interval
```

In some cases multiple merges will occur after a split so that the cluster may contain all available partitions. In this case, where multiple merges occur, time should be allowed for all of these to complete, and as there may be as many as three merges occurring sequentially the total delay should be no more

than the following:

10 \* *MERGE3.max\_interval*



### IMPORTANT

The amount of time taken in the formulas above is how long it takes JBoss Data Grid to install a cluster view without the leavers; however, as JBoss Data Grid runs inside a JVM excessive Garbage Collection (GC) times can increase this time beyond the failure detection outlined above. JBoss Data Grid has no control over these GC times, and excessive GC on the coordinator can delay this detection by an amount equal to the GC time.

In addition, when merging cluster views JBoss Data Grid tries to confirm all members are present; however, there is no upper bound on waiting for these responses, and merging the cluster views may be delayed due to networking issues.

[Report a bug](#)

## 32.4. DETECTING AND RECOVERING FROM SUCCESSIVE CRASHED NODES

Red Hat JBoss Data Grid is unable to distinguish whether a node left the cluster because of a process or machine crash, or because of a network failure.

If a single node exits the cluster, and if the value of *numOwners* is greater than 1, the cluster remains available and JBoss Data Grid attempts to create new replicas of the lost data. However, if additional nodes crash during this rebalancing process, it is possible that for some entries, all copies of its data have left the node and therefore cannot be recovered.

The recommended way to protect the data grid against successive crashed nodes is to enable partition handling (see [Section 32.6, “Configure Partition Handling”](#) for instructions) and to set an appropriately high value for *numOwners* to ensure that even if a large number of nodes leave the cluster in rapid succession, JBoss Data Grid is able to rebalance the nodes to recover the lost data.

[Report a bug](#)

## 32.5. NETWORK PARTITION RECOVERY EXAMPLES

The following examples illustrate how network partitions occur in Red Hat JBoss Data Grid clusters and how they are dealt with and eventually merged. The following examples scenarios are described in detail:

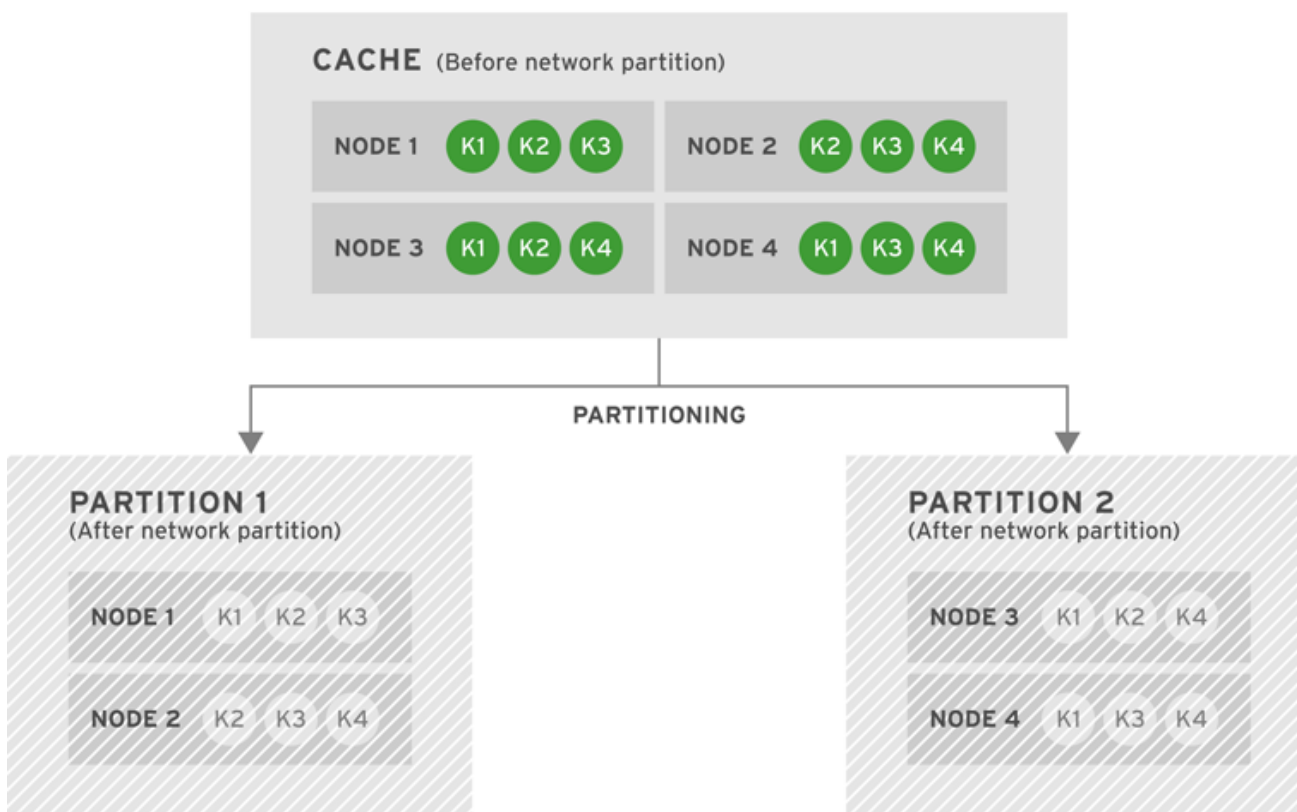
1. A distributed four node cluster with *numOwners* set to 3 at [Section 32.5.1, “Distributed 4-Node Cache Example With 3 NumOwners”](#)
2. A distributed four node cluster with *numOwners* set to 2 at [Section 32.5.2, “Distributed 4-Node Cache Example With 2 NumOwners”](#)
3. A distributed five node cluster with *numOwners* set to 3 at [Section 32.5.3, “Distributed 5-Node Cache Example With 3 NumOwners”](#)

4. A replicated four node cluster with *numOwners* set to 4 at [Section 32.5.4, “Replicated 4-Node Cache Example With 4 NumOwners”](#)
5. A replicated five node cluster with *numOwners* set to 5 at [Section 32.5.5, “Replicated 5-Node Cache Example With 5 NumOwners”](#)
6. A replicated eight node cluster with *numOwners* set to 8 at [Section 32.5.6, “Replicated 8-Node Cache Example With 8 NumOwners”](#)

[Report a bug](#)

### 32.5.1. Distributed 4-Node Cache Example With 3 NumOwners


The first example scenario includes a four-node distributed cache that contains four data entries (**k1**, **k2**, **k3**, and **k4**). For this cache, *numOwners* equals 3, which means that each data entry must have three copies on various nodes in the cache.



JBOSS\_11\_335370\_0415

**Figure 32.1. Cache Before and After a Network Partition**

As seen in the diagram, after the network partition occurs, Node 1 and Node 2 form Partition 1 while Node 3 and Node 4 form a Partition 2. After the split, the two partitions enter into Degraded Mode (represented by grayed-out nodes in the diagram) because neither has at least 3 (the value of *numOwners*) nodes left from the last stable view. As a result, none of the four entries (**k1**, **k2**, **k3**, and **k4**) are available for reads or writes. No new entries can be written in either degraded partition, as neither partition can store 3 copies of an entry.



JBoss\_12\_335370\_0415

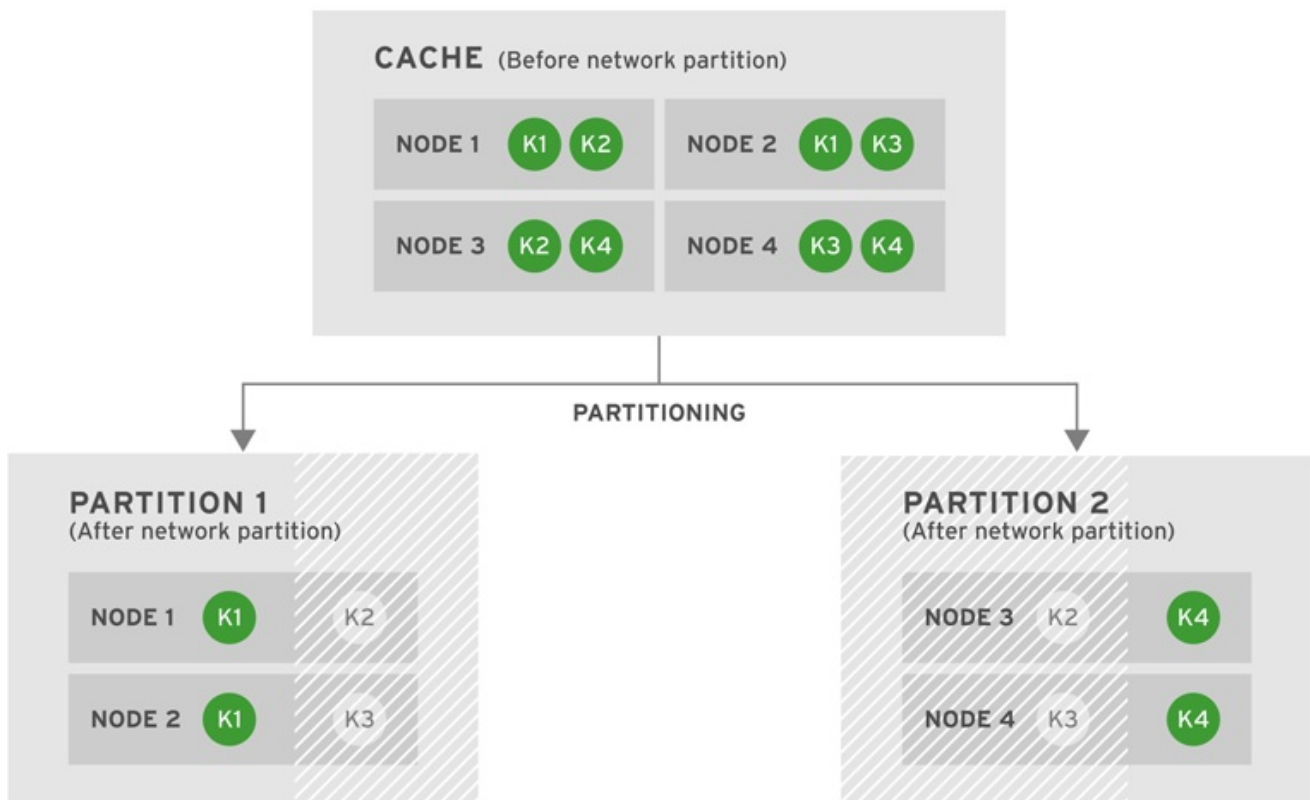
**Figure 32.2. Cache After Partitions Are Merged**

JBoss Data Grid subsequently merges the two split partitions. No state transfer is required and the new merged cache is subsequently in Available Mode with four nodes and four entries (**k1**, **k2**, **k3**, and **k4**).

[Report a bug](#)

### 32.5.2. Distributed 4-Node Cache Example With 2 NumOwners

The second example scenario includes a distributed cache with four nodes. In this scenario, *numOwners* equals 2, so the four data entries (**k1**, **k2**, **k3** and **k4**) have two copies each in the cache.

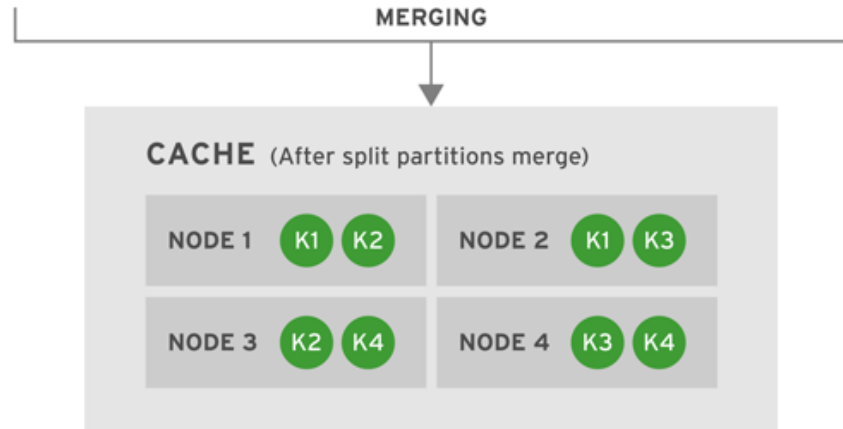


JBoss\_13\_335370\_0615

**Figure 32.3. Cache Before and After a Network Partition**

After the network partition occurs, Partitions 1 and 2 enter Degraded mode (depicted in the diagram as

grayed-out nodes). Within each partition, an entry will only be available for read or write operations if both its copies are in the same partition. In Partition 1, the data entry **k1** is available for reads and writes because *numOwners* equals 2 and both copies of the entry remain in Partition 1. In Partition 2, **k4** is available for reads and writes for the same reason. The entries **k2** and **k3** become unavailable in both partitions, as neither partition contains all copies of these entries. A new entry **k5** can be written to a partition only if that partition were to own both copies of **k5**.



JBOSSE\_14\_335370\_0415

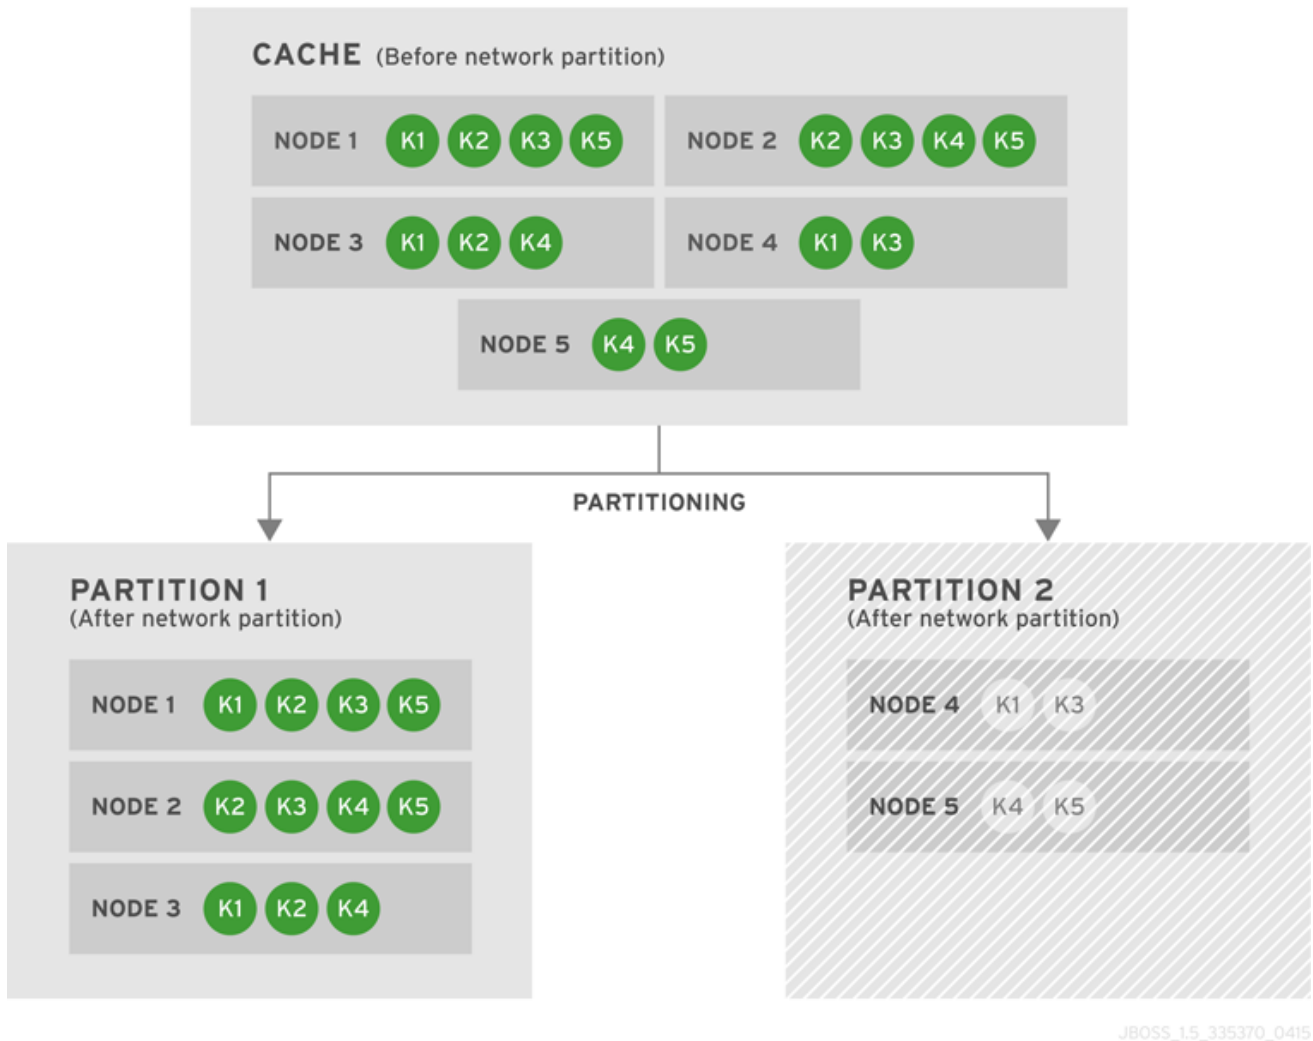
**Figure 32.4. Cache After Partitions Are Merged**

JBoss Data Grid subsequently merges the two split partitions into a single cache. No state transfer is required and the cache returns to Available Mode with four nodes and four data entries (**k1**, **k2**, **k3** and **k4**).

[Report a bug](#)

### 32.5.3. Distributed 5-Node Cache Example With 3 NumOwners

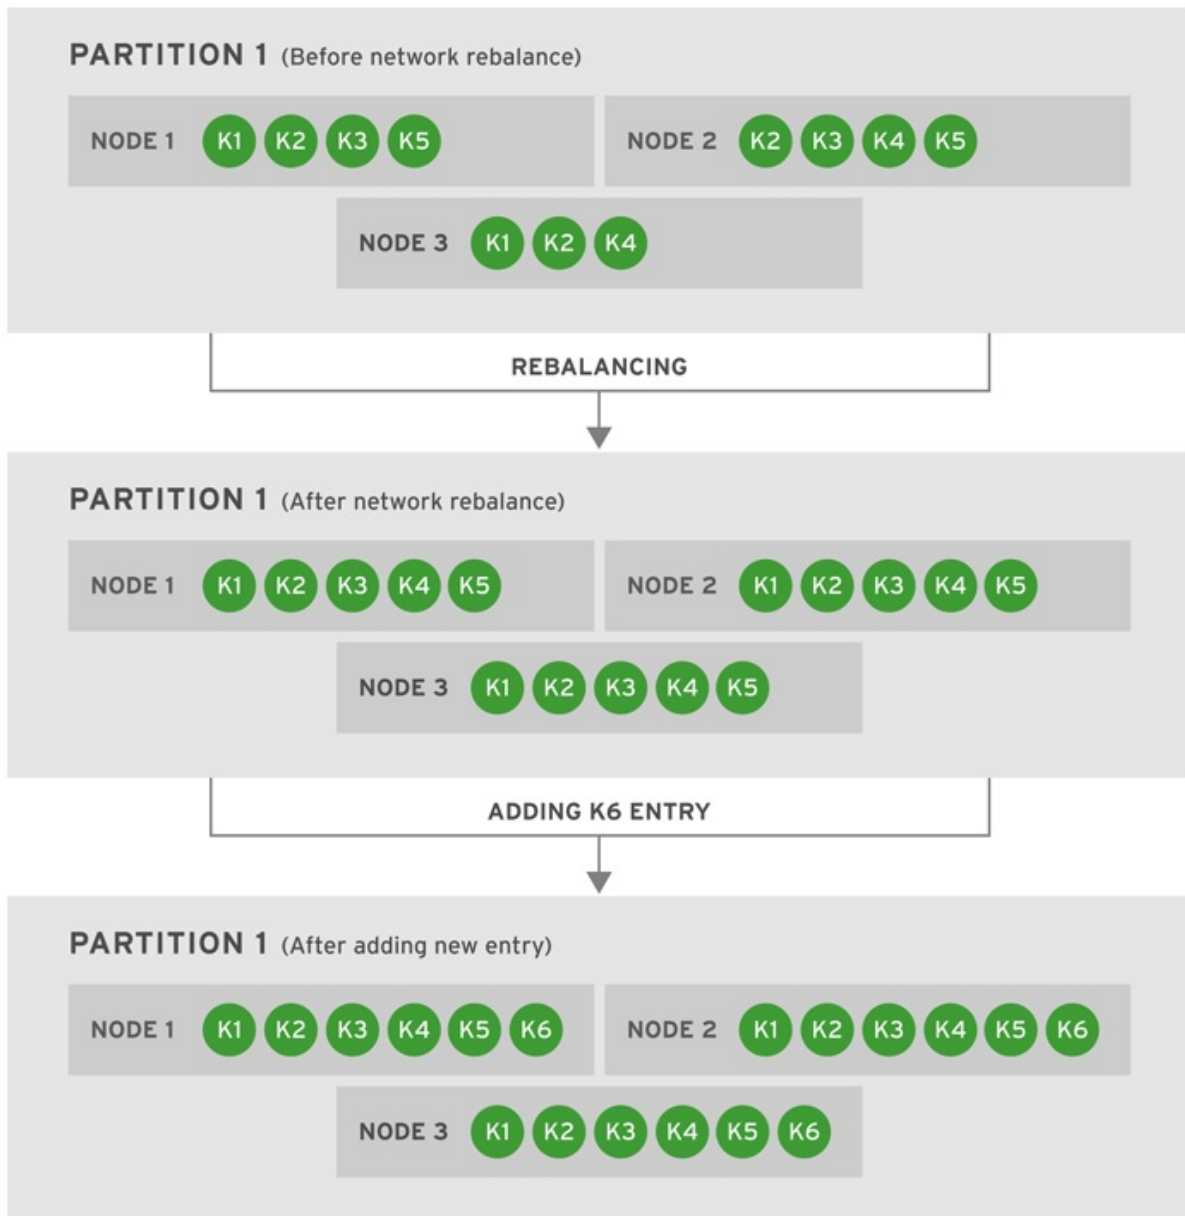
The third example scenario includes a distributed cache with five nodes and with *numOwners* equal to 3.



**Figure 32.5. Cache Before and After a Network Partition**

After the network partition occurs, the cache splits into two partitions. Partition 1 includes Node 1, Node 2, and Node 3 and Partition 2 includes Node 4 and Node 5. Partition 2 is Degraded because it does not include the majority of nodes from the total number of nodes in the cache. Partition 1 remains Available because it has the majority of nodes and lost less than *numOwners* nodes.

No new entries can be added to Partition 2 because this partition is Degraded and it cannot own all copies of the data.




JBOSS\_16\_335370\_0415

**Figure 32.6. Partition 1 Rebalances and Another Entry is Added**

After the partition split, Partition 1 retains the majority of nodes and therefore can rebalance itself by creating copies to replace the missing entries. As displayed in the diagram above, rebalancing ensures that there are three copies of each entry (*numOwners* equals 3) in the cache. As a result, each of the three nodes contains a copy of every entry in the cache. Next, we add a new entry, **k6**, to the cache. Since the *numOwners* value is still 3, and there are three nodes in Partition 1, each node includes a copy of **k6**.





JBoss\_17\_335370\_0415


**Figure 32.7. Cache After Partitions Are Merged**

Eventually, Partition 1 and 2 are merged into a cache. Since only three copies are required for each data entry (*numOwners*=3), JBoss Data Grid rebalances the nodes so that the data entries are distributed between the four nodes in the cache. The new combined cache becomes fully available.

[Report a bug](#)

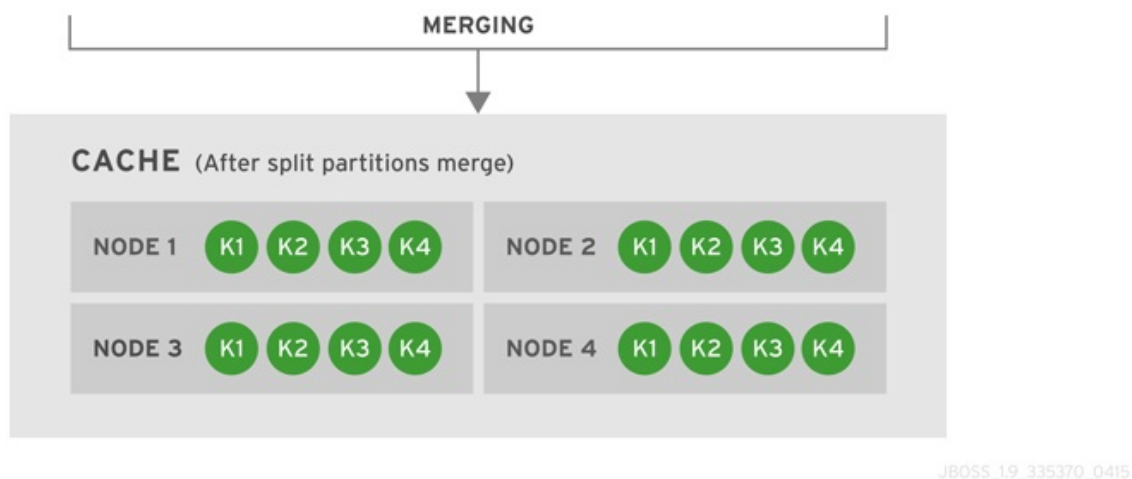
#### 32.5.4. Replicated 4-Node Cache Example With 4 NumOwners

The fourth example scenario includes a replicated cache with four nodes and with *numOwners* equal to 4.



**Figure 32.8. Cache Before and After a Network Partition**

After a network partition occurs, Partition 1 contains Node 1 and Node 2 while Node 3 and Node 4 are in Partition 2. Both partitions enter Degraded Mode because neither has a simple majority of nodes. All four keys (**k1**, **k2**, **k3**, and **k4**) are unavailable for reads and writes because neither of the two partitions owns all copies of any of the four keys.



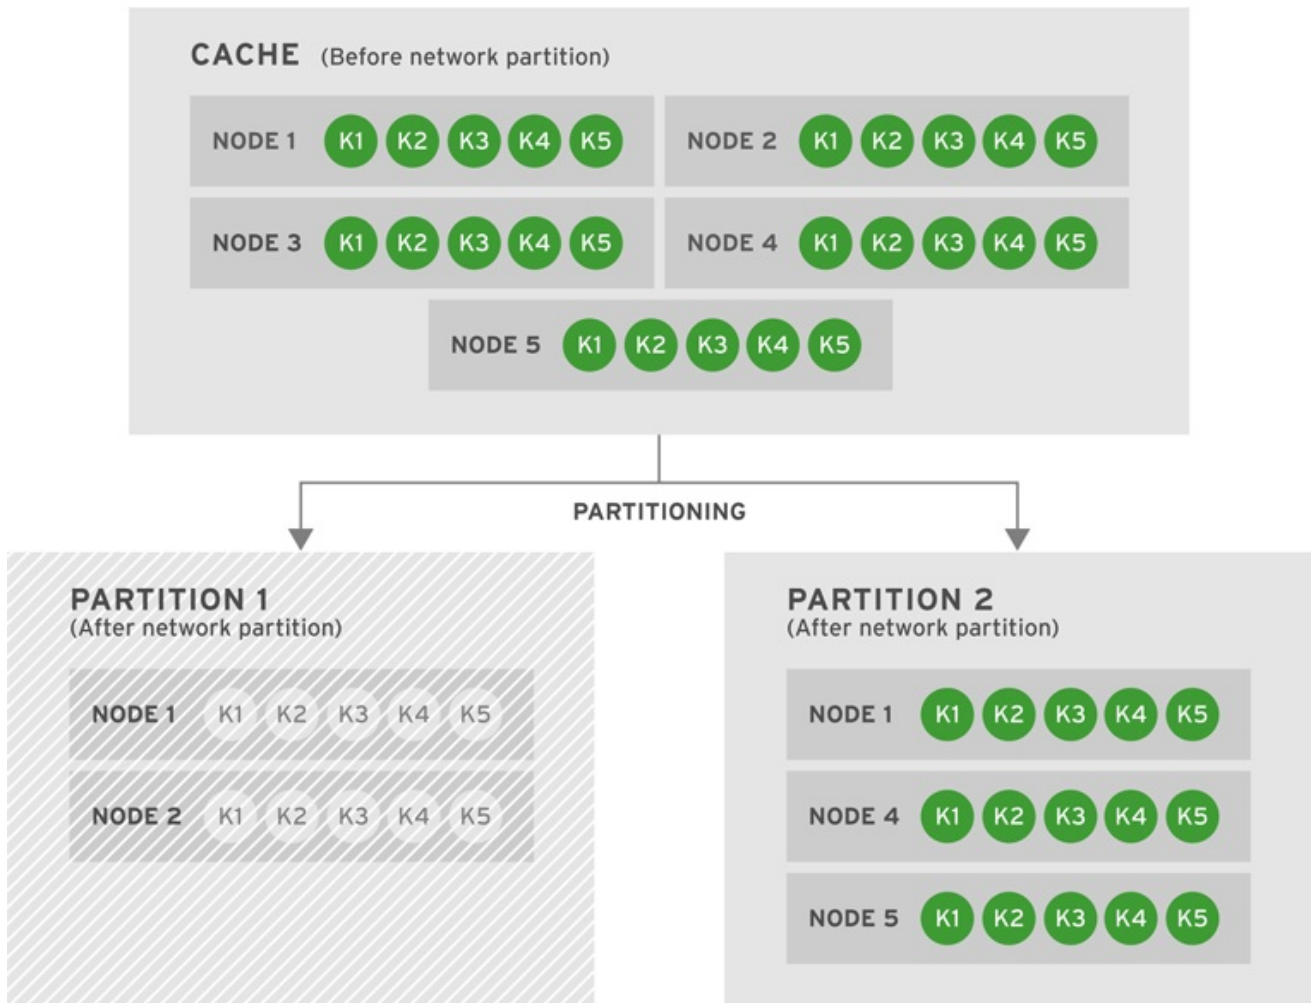
**Figure 32.9. Cache After Partitions Are Merged**

JBoss Data Grid subsequently merges the two split partitions into a single cache. No state transfer is required and the cache returns to its original state in Available Mode with four nodes and four data entries (**k1**, **k2**, **k3**, and **k4**).

[Report a bug](#)

### 32.5.5. Replicated 5-Node Cache Example With 5 NumOwners


The fifth example scenario includes a replicated cache with five nodes and with *numOwners* equal to 5.



JBOSS\_110\_335370\_0415

**Figure 32.10. Cache Before and After a Network Partition**

After a network partition occurs, the cache splits into two partitions. Partition 1 contains Node 1 and Node 2 and Partition 2 includes Node 3, Node 4, and Node 5. Partition 1 enters Degraded Mode (indicated by the grayed-out nodes) because it does not contain the majority of nodes. Partition 2, however, remains available.



JBOSS\_1.11\_335370\_0415

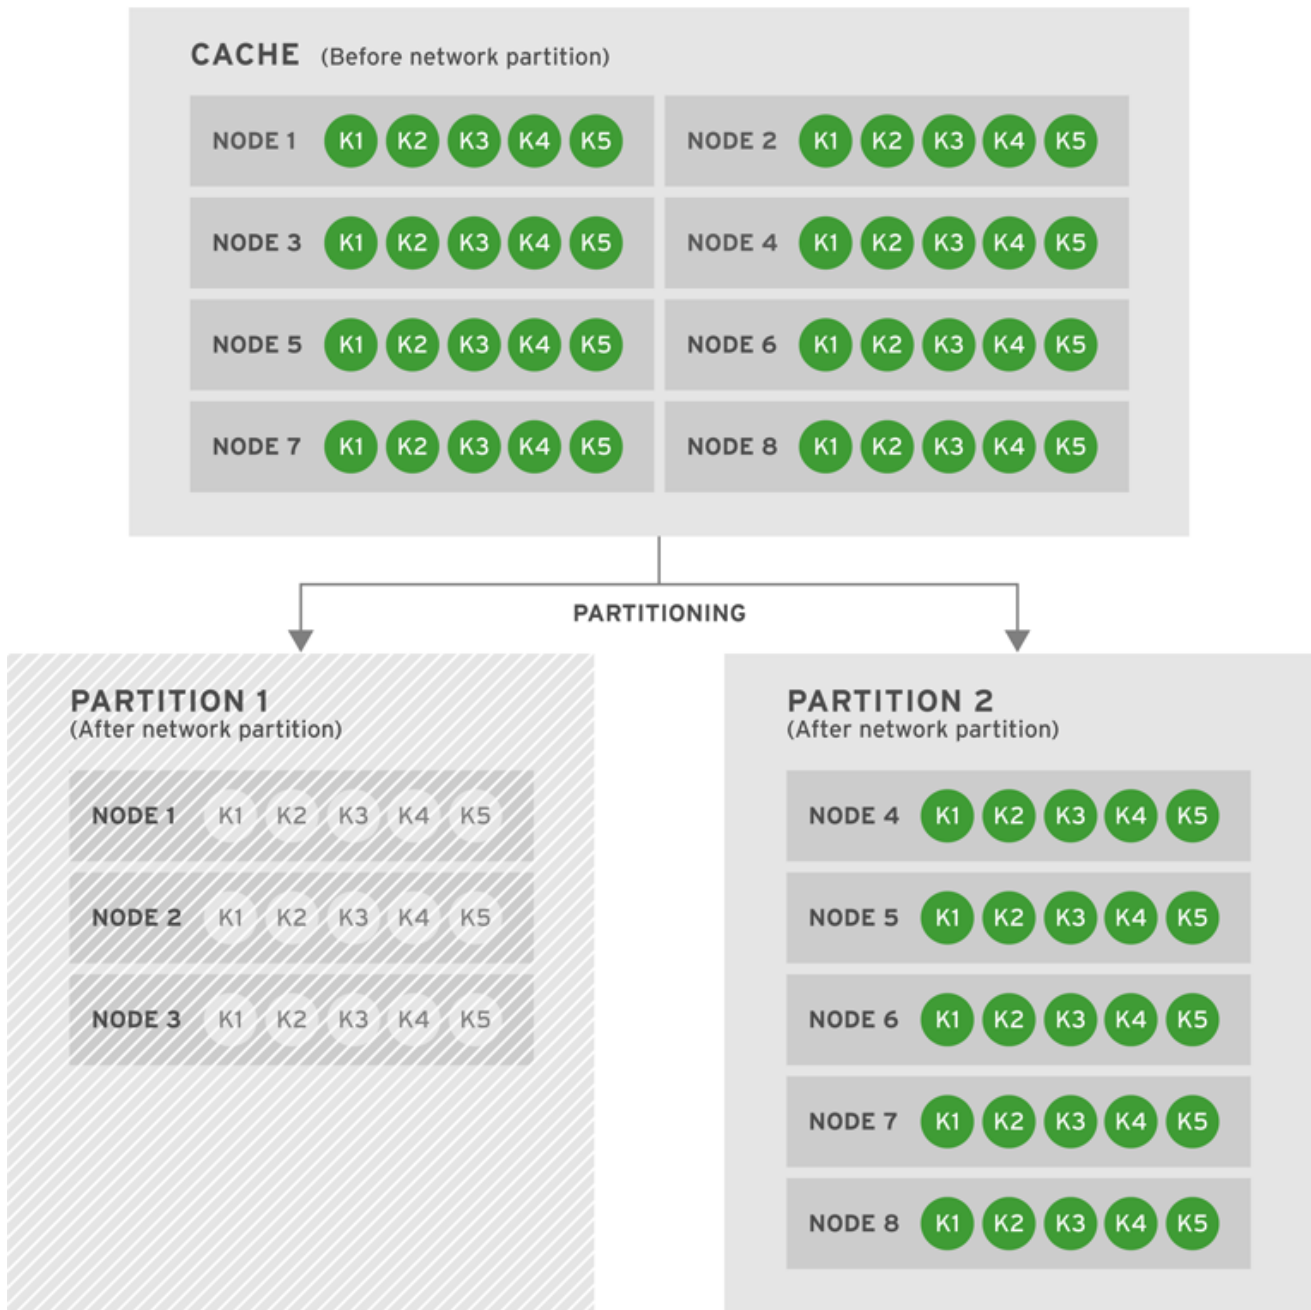
**Figure 32.11. Both Partitions Are Merged Into One Cache**

When JBoss Data Grid merges partitions in this example, Partition 2, which was fully available, is considered the primary partition. State is transferred from Partition 1 and to Partition 2. The combined cache becomes fully available."

[Report a bug](#)

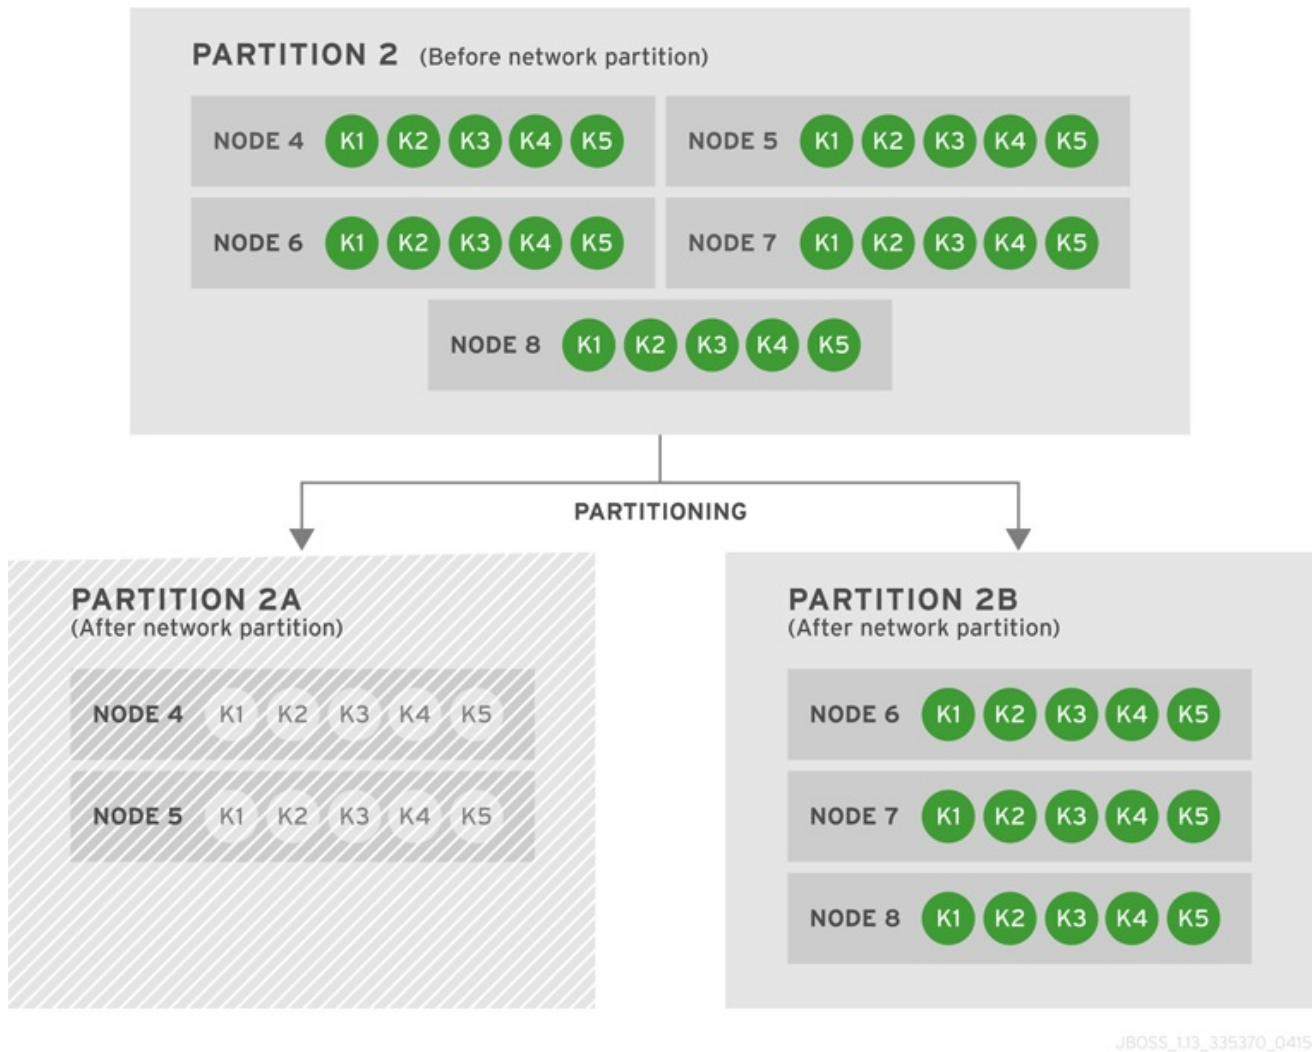
### 32.5.6. Replicated 8-Node Cache Example With 8 NumOwners

The sixth scenario is for a replicated cache with eight nodes and *numOwners* equal to 8.



**Figure 32.12. Cache Before and After a Network Partition**

A network partition splits the cluster into Partition 1 with 3 nodes and Partition 2 with 5 nodes. Partition 1 enters Degraded state, but Partition 2 remains Available.




**Figure 32.13. Partition 2 Further Splits into Partitions 2A and 2B**

Now another network partition affects Partition 2, which subsequently splits further into Partition 2A and 2B. Partition 2A contains Node 4 and Node 5 while Partition 2B contains Node 6, Node 7, and Node 8. Partition 2A enters Degraded Mode because it does not contain the majority of nodes. However, Partition 2B remains Available.

### Potential Resolution Scenarios


There are four potential resolutions for the caches from this scenario:

- Case 1: Partitions 2A and 2B Merge
- Case 2: Partition 1 and 2A Merge
- Case 3: Partition 1 and 2B Merge
- Case 4: Partition 1, Partition 2A, and Partition 2B Merge Together



**Figure 32.14. Case 1: Partitions 2A and 2B Merge**

The first potential resolution to the partitioned network involves Partition 2B's state information being copied into Partition 2A. The result is Partition 2, which contains Node 5, Node 6, Node 7, and Node 8. The newly merged partition becomes Available.




JBOSS\_115\_335370\_0415

**Figure 32.15. Case 2: Partition 1 and 2A Merge**

The second potential resolution to the partitioned network involves Partition 1 and Partition 2A merging. The combined partition contains Node 1, Node 2, Node 3, Node 4, and Node 5. As neither partition has the latest stable topology, the resulting merged partition remains in Degraded mode.

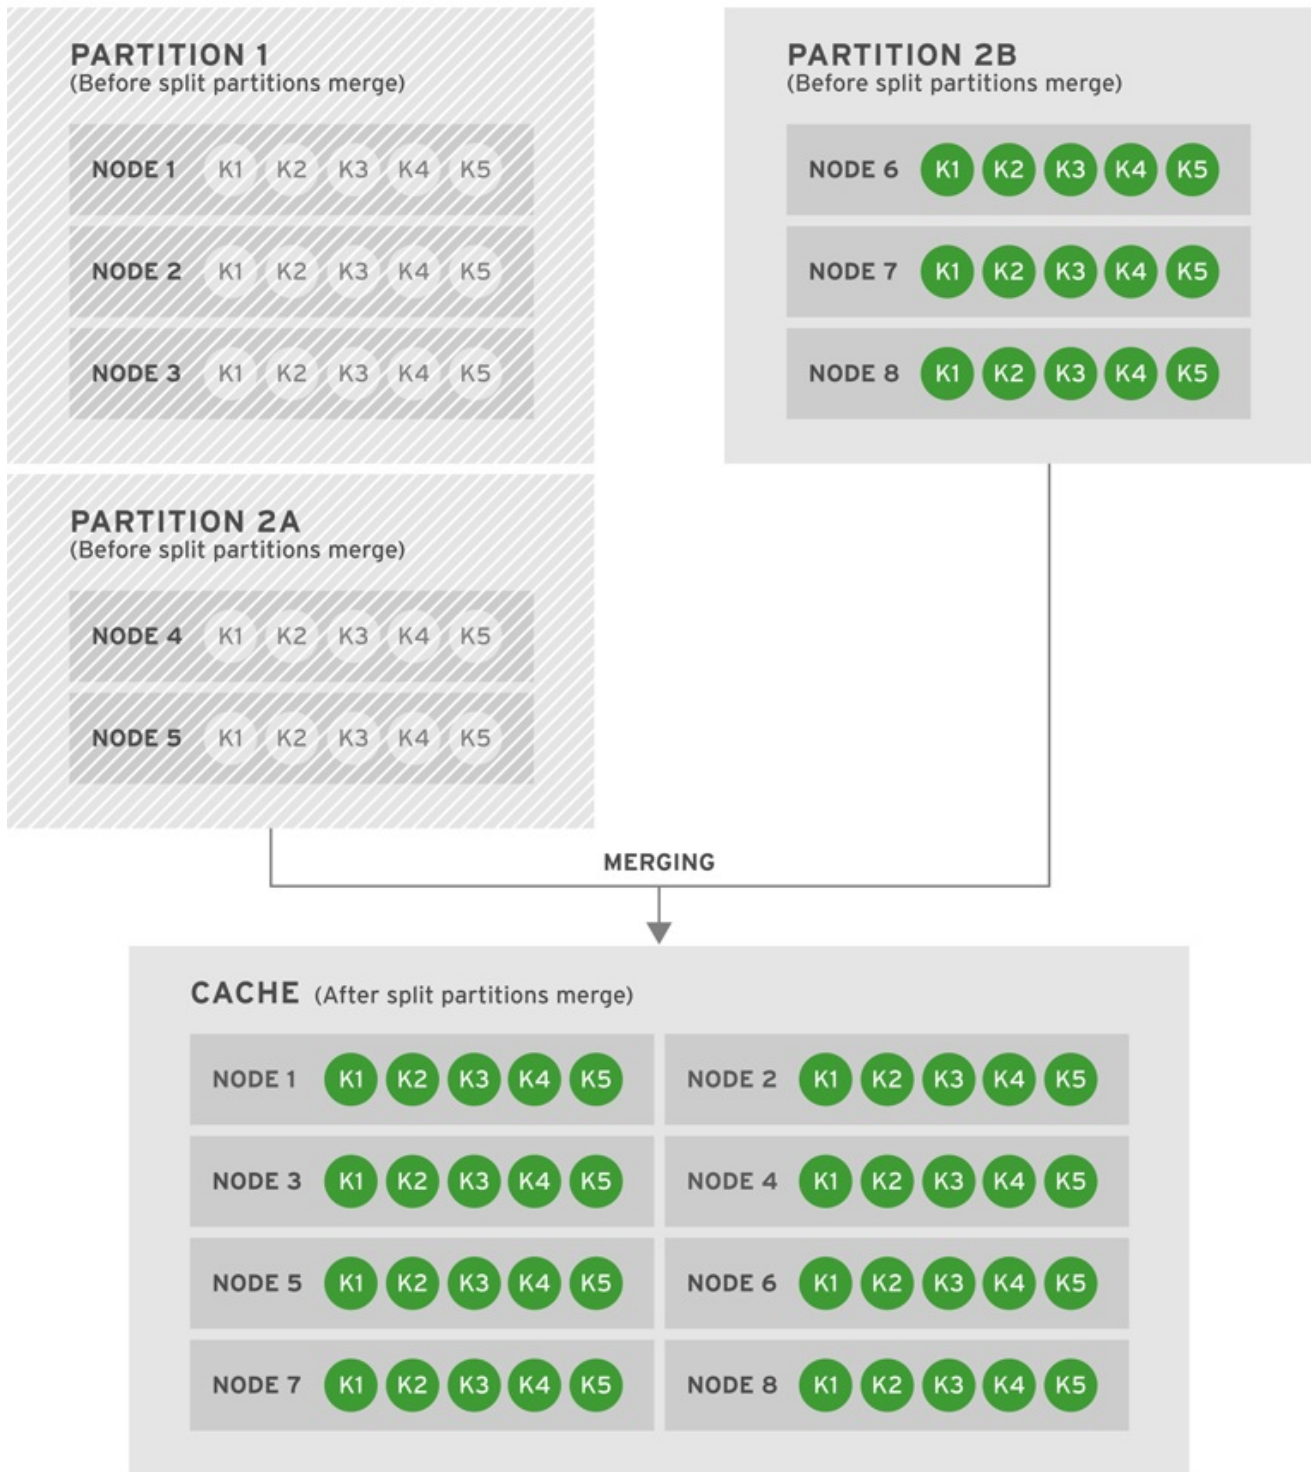




JBOSS\_116\_335370\_0415

**Figure 32.16. Case 3: Partition 1 and 2B Merge**

The third potential resolution to the partitioned network involves Partition 1 and Partition 2B merging. Partition 1 receives state information from Partition 2B, and the combined partition becomes Available.



JBOSS\_117\_335370\_0415

**Figure 32.17. Case 4: Partition 1, Partition 2A, and Partition 2B Merge Together**

The fourth and final potential resolution to the partitioned network involves Partition 1, Partition 2A, and Partition 2B merging to form Partition 1. The state is transferred from Partition 2B to both partitions 1 and 2A. The resulting cache contains eight nodes (Node 1, Node 2, Node 3, Node 4, Node 5, Node 6, Node 7, and Node 8) and is Available.

[Report a bug](#)

## 32.6. CONFIGURE PARTITION HANDLING

In Red Hat JBoss Data Grid, partition handling is disabled as a default.

### Declarative Configuration (Library Mode)

Enable partition handling declaratively as follows:

```
<namedCache name="myCache">
  <clustering mode="distribution">
    <partitionHandling enabled="true" />
  </clustering>
</namedCache>
```

### Programmatic Configuration (Library Mode)

Enable partition handling programmatically as follows:

```
ConfigurationBuilder dcc = new ConfigurationBuilder();
dcc.clustering().partitionHandling().enabled(true);
```

### Declarative Configuration (Remote Client-server Mode)

Enable partition handling declaratively in remote client-server mode by using the following configuration:

```
<subsystem xmlns="urn:infinispan:server:core:6.2" default-cache-
container="clustered">
  <cache-container name="clustered" default-cache="default"
statistics="true">
    <distributed-cache name="default" mode="SYNC" segments="20"
owners="2"
                    remote-timeout="30000" start="EAGER">
      <partition-handling enabled="true" />
      <locking isolation="READ_COMMITTED" acquire-timeout="30000"
concurrency-level="1000" striping="false"/>
      <transaction mode="NONE"/>
    </distributed-cache>
  </cache-container>
</subsystem>
```

[Report a bug](#)

## APPENDIX A. RECOMMENDED JGROUPS VALUES FOR JBOSS DATA GRID

### A.1. SUPPORTED JGROUPS PROTOCOLS

The table contains a list of the JGroups protocols supported in JBoss Data Grid.

**Table A.1. Supported JGroups Protocols**

Protocol	Details
TCP	<p>TCP/IP is a replacement transport for UDP in situations where IP multicast cannot be used, such as operations over a WAN where routers may discard IP multicast packets.</p> <p>TCP is a transport protocol used to send unicast and multicast messages.</p> <ul style="list-style-type: none"> <li>• When sending multicast messages, TCP sends multiple unicast messages.</li> <li>• When using TCP, each message to all cluster members is sent as multiple unicast messages, or one to each member.</li> </ul> <p>As IP multicasting cannot be used to discover initial members, another mechanism must be used to find initial membership.</p> <p>Red Hat JBoss Data Grid's Hot Rod is a custom TCP client/server protocol.</p>
UDP	<p>UDP is a transport protocol that uses:</p> <ul style="list-style-type: none"> <li>• IP multicasting to send messages to all members of a cluster.</li> <li>• UDP datagrams for unicast messages, which are sent to a single member.</li> </ul> <p>When the UDP transport is started, it opens a unicast socket and a multicast socket. The unicast socket is used to send and receive unicast messages, the multicast socket sends and receives multicast sockets. The physical address of the channel will be the same as the address and port number of the unicast socket.</p>

Protocol	Details
PING	<p>The PING protocol is used for the initial discovery of members. It is used to detect the coordinator, which is the oldest member, by multicasting PING requests to an IP multicast address.</p> <p>Each member responds to the ping with a packet containing the coordinator's and their own address. After a specified number of milliseconds (N) or replies (M), the joiner determines the coordinator from the responses and sends it a JOIN request (handled by GMS). If there is no response, the joiner is considered the first member of the group.</p> <p>PING differs from TCPPING because it used dynamic discovery, which means that a member does not need to know in advance where the other cluster members are. PING uses the transport's IP multicasting abilities to send a discovery request to the cluster. As a result, PING requires UDP as transport.</p>
TCPPING	<p>The TCPPING protocol uses a set of known members and pings them for discovery. This protocol has a static configuration.</p>
MPING	<p>The MPING (Multicast PING) protocol uses IP multicast to discover the initial membership. It can be used with all transports, but is usually used in combination with TCP.</p>
S3_PING	<p>S3_PING is a discovery protocol that is ideal for use with Amazon's Elastic Compute Cloud (EC2) because EC2 does not allow multicast and therefore MPING is not allowed.</p> <p>Each EC2 instance adds a small file to an S3 data container, known as a bucket. Each instance then reads the files in the bucket to discover the other members of the cluster.</p>
MERGE3	<p>The MERGE3 protocol is available in JGroups 3.1 onwards. Unlike MERGE2, in MERGE3, all members periodically send an INFO message with their address (UUID), logical name, physical address and View ID. Periodically, each coordinator reviews the INFO details to ensure that there are no inconsistencies.</p>
FD_ALL	<p>Used for failure detection, FD_ALL uses a simple heartbeat protocol. Each member maintains a table of all other members (except itself) and periodically multicasts a heartbeat. For example, when data or a heartbeat from P is received, the timestamp for P is set to the current time. Periodically, expired members are identified using the timestamp values.</p>

Protocol	Details
FD_SOCK	<p>FD_SOCK is a failure detection protocol based on a ring of TCP sockets created between cluster members. Each cluster member connects to its neighbor (the last member connects to the first member), which forms a ring. Member B is suspected when its neighbor A detects an abnormal closing of its TCP socket (usually due to node B crashing). However, if member B is leaving gracefully, it informs member A and does not become suspected when it does exit.</p>
FD_HOST	<p>FD_HOST is a failure detection protocol that detects the crashing or hanging of entire hosts and suspects all cluster members of that host through ICMP ping messages or custom commands. FD_HOST does not detect the crashing or hanging of single members on the local hosts, but only checks whether all the other hosts in the cluster are live and available. It is therefore used in conjunction with other failure detection protocols such as FD_ALL and FD_SOCK. This protocol is typically used when multiple cluster members are running on the same physical box.</p> <p>The FD_HOST protocol is supported on Windows for JBoss Data Grid. The <code>cmd</code> parameter must be set to <code>ping.exe</code> and the ping count must be specified.</p>
VERIFY_SUSPECT	<p>The VERIFY_SUSPECT protocol verifies whether a suspected member is dead by pinging the member before excluding it. If the member responds, the suspect message is discarded.</p>
NAKACK2	<p>The NAKACK2 protocol is a successor to the NAKACK protocol and was introduced in JGroups 3.1.</p> <p>The NACKACK2 protocol is used for multicast messages and uses NAK. Each message is tagged with a sequence number. The receiver tracks the sequence numbers and delivers the messages in order. When a gap in the sequence numbers is detected, the receiver asks the sender to retransmit the missing message.</p>

Protocol	Details
UNICAST3	<p>The UNICAST3 protocol provides reliable delivery (no message sent by a sender is lost because they are sent in a numbered sequence) and uses the FIFO (First In First Out) properties for point to point messages between a sender and a receiver.</p> <p>UNICAST3 uses positive acks for retransmission. For example, sender A keeps sending message M until receiver B receives message M and returns an ack to indicate a successful delivery. Sender A keeps resending message M until it receives an ack from B, until B leaves the cluster, or A crashes.</p>
STABLE	<p>The STABLE protocol is a garbage collector for messages that have been viewed by all members in a cluster. Each member stores all messages because retransmission may be required. A message can only be removed from the retransmission buffers when all members have seen the message. The STABLE protocol periodically gossips its highest and lowest messages seen. The lowest value is used to compute the min (all lowest sequence numbers for all members) and messages with a sequence number below the min value can be discarded</p>
GMS	<p>The GMS protocol is the group membership protocol. This protocol handles joins/leaves/crashes (suspicions) and emits new views accordingly.</p>
MFC	<p>MFC is the Multicast version of the flow control protocol.</p>
UFC	<p>UFC is the Unicast version of the flow control protocol.</p>
FRAG2	<p>The FRAG2 protocol fragments large messages into smaller ones and then sends the smaller messages. At the receiver side, the smaller fragments are reassembled into larger, complete messages and delivered to the application. FRAG2 is used for both multicast and unicast messages.</p>


Protocol	Details
ENCRYPT	<p>JGroups includes the ENCRYPT protocol to provide encryption for cluster traffic. By default, encryption only encrypts the message body; it does not encrypt message headers. To encrypt the entire message, including all headers, as well as destination and source addresses, the property <code>encrypt_entire_message</code> must be true. ENCRYPT must also be below any protocols with headers that must be encrypted.</p> <p>The ENCRYPT layer is used to encrypt and decrypt communication in JGroups. The JGroups ENCRYPT protocol can be used in two ways:</p> <ul style="list-style-type: none"> <li>• Configured with a <code>secretKey</code> in a keystore.</li> <li>• Configured with algorithms and key sizes.</li> </ul> <p>Each message is identified as encrypted with a specific encryption header identifying the encrypt header and an MD5 digest identifying the version of the key being used to encrypt and decrypt messages.</p>
SASL	<p>The SASL (Simple Authentication and Security Layer) protocol is a framework that provides authentication and data security services in connection-oriented protocols using replaceable mechanisms. Additionally, SASL provides a structured interface between protocols and mechanisms.</p>
RELAY2	<p>The RELAY protocol bridges two remote clusters by creating a connection between one node in each site. This allows multicast messages sent out in one site to be relayed to the other and vice versa.</p> <p>JGroups includes the RELAY2 protocol, which is used for communication between sites in Red Hat JBoss Data Grid's Cross-Site Replication.</p> <p>The RELAY2 protocol works similarly to RELAY but with slight differences. Unlike RELAY, the RELAY2 protocol:</p> <ul style="list-style-type: none"> <li>• connects more than two sites.</li> <li>• connects sites that operate autonomously and are unaware of each other.</li> <li>• offers both unicast and multicast routing between sites.</li> </ul>

[Report a bug](#)

## A.2. TCP DEFAULT AND RECOMMENDED VALUES

To learn more about JGroups and using TCP and UDP, see [Section 26.2, “Configure JGroups \(Library Mode\)”](#).



**NOTE**

- Values in **JGroups Default Value** indicate values that are configured internally to JGroups, but may be overridden by a custom configuration file or by a JGroups configuration file shipped with JBoss Data Grid.
- Values in **JBoss Data Grid Configured Values** indicate values that are in use by default when using one of the configuration files for JGroups as shipped with JBoss Data Grid. It is recommended to use these values when custom configuration files for JGroups are in use with JBoss Data Grid.

For more information on the configuration files included with JBoss Data Grid refer to [Section A.4, “JBoss Data Grid JGroups Configuration Files”](#).

**Table A.2. Recommended and Default Values for TCP**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
bind_addr	Any non-loopback	Set address on specific interface
bind_port	Any free port	Set specific port
loopback	true	Same as default
port_range	50	Set based on desired range of ports
recv_buf_size	150,000	Same as default
send_buf_size	150,000	640,000
use_send_queues	true	Same as default
sock_conn_timeout	2,000	300
max_bundle_size	64,000	64,000
enable_diagnostics	true	false
thread_pool.enabled	true	Same as default
thread_pool.min_threads	2	This should equal the number of nodes

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
thread_pool.max_threads	30	This should be higher than <b><i>thread_pool.min_threads</i></b> . For example, for a smaller grid (2-10 nodes), set this value to twice the number of nodes, but for a larger grid (20 or more nodes), the ratio should be lower. As an example, if a grid contains 20 nodes, set this value to 25 and if the grid contains 100 nodes, set the value to 110.
thread_pool.keep_alive_time	30,000	60,000
thread_pool.queue_enabled	true	false
thread_pool.queue_max_size	500	None, queue should be disabled
thread_pool.rejection_policy	Discard	Same as default
internal_thread_pool.enabled	true	Same as default
internal_thread_pool.min_threads	2	5
internal_thread_pool.max_threads	4	20
internal_thread_pool.keep_alive_time	30,000	60,000
internal_thread_pool.queue_enabled	true	false
internal_thread_pool.rejection_policy	Discard	Abort
oob_thread_pool.enabled	true	Same as default
oob_thread_pool.min_threads	2	20 or higher
oob_thread_pool.max_threads	10	200 or higher, depending on the load
oob_thread_pool.keep_alive_time	30,000	60,000
oob_thread_pool.queue_enabled	true	false

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
<code>oob_thread_pool.queue_max_size</code>	500	None, queue should be disabled
<code>oob_thread_pool.rejection_policy</code>	Discard	Same as default

**NOTE**

Red Hat JBoss Data Grid 6.5 uses JGroups 3.6.3, in which the TCPPING timeout value is removed and the ***pbcast.GMS join\_timeout*** value indicates the timeout period instead.

**Recommended Values for S3\_PING**

See [Section 27.2, “S3\\_PING Configuration Options”](#) for details about configuring S3\_PING for JBoss Data Grid.

**Table A.3. Recommended Values for MPING**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
<code>bind_addr</code>	Any non-loopback	Set address on specific interface
<code>break_on_coord_rsp</code>	true	Same as default
<code>mcast_addr</code>	230.5.6.7	Same as default
<code>mcast_port</code>	7555	Same as default
<code>ip_ttl</code>	8	2

**NOTE**

In JGroups 3.6.1, the MPING timeout value is removed and the ***pbcast.GMS join\_timeout*** value indicates the timeout period instead.

**Table A.4. Recommended Values for MERGE3**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
<code>min_interval</code>	1,000	10,000
<code>max_interval</code>	10,000	30,000

Table A.5. Recommended Values for FD SOCK

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
client_bind_por	0 (randomly selects a port and uses it)	Same as default
get_cache_timeout	1000 milliseconds	Same as default
keep_alive	true	Same as default
num_tries	3	Same as default
start_port	0 (randomly selects a port and uses it)	Same as default
suspect_msg_interval	5000 milliseconds.	Same as default

Table A.6. Recommended Values for FD\_ALL

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
timeout	40,000	60,000. The FD_ALL timeout value is set to two times the longest possible stop the world garbage collection pause in the CMS garbage collector. In a well-tuned JVM, the longest pause is proportional to heap size and should not exceed 1 second per GB of heap. For example, an 8GB heap should not have a pause longer than 8 seconds, so the FD_ALL timeout value must be set to 16 seconds. If longer garbage collection pauses are used, then this timeout value should be increased to avoid false failure detection on a node.
interval	8,000	15,000. The FD_ALL <i>interval</i> value must be at least four times smaller than the value set for FD_ALL's <i>timeout</i> value.
timeout_check_interval	2,000	5,000

Table A.7. Recommended Values for FD\_HOST

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
check_timeout	3,000	5,000
cmd	InetAddress.isReachable() (ICMP ping)	-
interval	20,000	15,000. The <b>interval</b> value for FD_HOST must be four times smaller than FD_HOST's <b>timeout</b> value.
timeout	60,000	60,000.

**Table A.8. Recommended Values for VERIFY\_SUSPECT**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
timeout	2,000	5,000

**Table A.9. Recommended Values for pbcast.NAKACK2**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
use_mcast_xmit	true	false
xmit_interval	1,000	Same as default
xmit_table_num_rows	50	50
xmit_table_msgs_per_row	10,000	1,024
xmit_table_max_compaction_time	10,000	30,000
max_msg_batch_size	100	Same as default
resend_last_seqno	false	true

**Table A.10. Recommended Values for UNICAST3**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
xmit_interval	500	Same as default

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
xmit_table_num_rows	100	50
xmit_table_msgs_per_row	1,0000	1,024
xmit_table_max_compaction_time	600,000	30,000
max_msg_batch_size	500	100
conn_close_timeout	60,000	No recommended value.
conn_expiry_timeout	120,000	0

**Table A.11. Recommended Values for pbcaster.STABLE**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
stability_delay	6,000	500
desired_avg_gossip	20,000	5,000
max_bytes	2,000,000	1,000,000

**Table A.12. Recommended Values for pbcaster.GMS**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
print_local_addr	true	false
join_timeout	5,000	15,000
view_bundling	true	Same as default

**Table A.13. Recommended Values for MFC**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
max_credits	500,000	2,000,000
min_threshold	0.40	Same as default

**Table A.14. Recommended Values for FRAG2**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
frag_size	60,000	Same as default

**Table A.15. Recommended Values for ENCRYPT**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
asymAlgorithm	RSA	-
asymInit	512	-
asymProvider	Bouncy Castle Provider	-
changeKeysOnViewChange	false	-
symAlgorithm	AES	-
symInit	128	-
symProvider	Bouncy Castle Provider	-

**NOTE**

See the *Red Hat JBoss Data Grid Developer Guide's JGroups ENCRYPT* section for details.

**Recommended Values for SASL**

See the *Red Hat JBoss Data Grid Developer Guide's User Authentication over Hot Rod Using SASL* section for details.

**Recommended Values for RELAY2**

See [Chapter 30, Set Up Cross-Datacenter Replication](#) for details.

[Report a bug](#)

**A.3. UDP DEFAULT AND RECOMMENDED VALUES**

To learn more about JGroups and using TCP and UDP, see [Section 26.2, "Configure JGroups \(Library Mode\)"](#).

**NOTE**

- Values in **JGroups Default Value** indicate values that are configured internally to JGroups, but may be overridden by a custom configuration file or by a JGroups configuration file shipped with JBoss Data Grid.
- Values in **JBoss Data Grid Configured Values** indicate values that are in use by default when using one of the configuration files for JGroups as shipped with JBoss Data Grid. It is recommended to use these values when custom configuration files for JGroups are in use with JBoss Data Grid.

For more information on the configuration files included with JBoss Data Grid refer to [Section A.4, “JBoss Data Grid JGroups Configuration Files”](#).


**Table A.16. Recommended Values for UDP**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
bind_addr	Any non-loopback	Set address on specific interface
bind_port	Any free port	Set specific port
loopback	true	true
port_range	50	Set based on desired range of ports
mcast_addr	228.8.8.8	Same as default
mcast_port	7600	Same as default
tos	8	Same as default
ucast_rcv_buf_size	64,000	20,000,000
ucast_send_buf_size	100,000	1,000,000
mcast_rcv_buf_size	500,000	25,000,000
mcast_send_buf_size	100,000	1,000,000
ip_ttl	8	2
thread_naming_pattern	cl	pl
max_bundle_size	64,000	Same as default
enable_diagnostics	true	false



Parameter	JGroups Default Value	JBoss Data Grid Configured Values
thread_pool.enabled	true	Same as default
thread_pool.min_threads	2	This should equal the number of nodes.
thread_pool.max_threads	30	This should be higher than thread_pool.min_threads. For example, for a smaller grid (2-10 nodes), set this value to twice the number of nodes, but for a larger grid (20 or more nodes), the ratio should be lower. As an example, if a grid contains 20 nodes, set this value to 25 and if the grid contains 100 nodes, set the value to 110.
thread_pool.keep_alive_time	30,000	60,000
thread_pool.queue_enabled	true	false
thread_pool.queue_max_size	500	None, queue should be disabled
thread_pool.rejection_policy	Discard	Same as default
internal_thread_pool.enabled	true	Same as default
internal_thread_pool.min_threads	2	5
internal_thread_pool.max_threads	4	20
internal_thread_pool.keep_alive_time	30,000	60,000
internal_thread_pool.queue_enabled	true	false
internal_thread_pool.rejection_policy	Discard	Abort
oob_thread_pool.enabled	true	Same as default
oob_thread_pool.min_threads	2	20 or higher
oob_thread_pool.max_threads	10	200 or higher based on the load
oob_thread_pool.keep_alive_time	30,000	60,000

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
<code>oob_thread_pool.queue_enabled</code>	true	false
<code>oob_thread_pool.queue_max_size</code>	500	None, queue should be disabled
<code>oob_thread_pool.rejection_policy</code>	Discard	Same as default

**NOTE**

In JGroups 3.5, the PING timeout value is removed and the `pbcast.GMS` *join\_timeout* value indicates the timeout period instead.

**Table A.17. Recommended Values for MERGE3**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
<code>min_interval</code>	1,000	10,000
<code>max_interval</code>	10,000	30,000

**Table A.18. Recommended Values for FD SOCK**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
<code>client_bind_port</code>	0 (randomly selects a port and uses it)	Same as default
<code>get_cache_timeout</code>	1000 milliseconds	Same as default
<code>keep_alive</code>	true	Same as default
<code>num_tries</code>	3	Same as default
<code>start_port</code>	0 (randomly selects a port and uses it)	Same as default
<code>suspect_msg_interval</code>	5000 milliseconds.	Same as default

**Table A.19. Recommended Values for FD\_ALL**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
timeout	40,000	60,000. The FD_ALL timeout value is set to two times the longest possible stop the world garbage collection pause in the CMS garbage collector. In a well-tuned JVM, the longest pause is proportional to heap size and should not exceed 1 second per GB of heap. For example, an 8GB heap should not have a pause longer than 8 seconds, so the FD_ALL timeout value must be set to 16 seconds. If longer garbage collection pauses are used, then this timeout value should be increased to avoid false failure detection on a node.
interval	8,000	15,000. The FD_ALL <i>interval</i> value must be at least four times smaller than the value set for FD_ALL's <i>timeout</i> value.
timeout_check_interval	2,000	5,000

Table A.20. Recommended Values for FD\_HOST

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
check_timeout	3,000	5,000
cmd	InetAddress.isReachable() (ICMP ping)	-
interval	20,000	15,000. The <i>interval</i> value for FD_HOST must be four times smaller than FD_HOST's <i>timeout</i> value.
timeout	-	60,000.

Table A.21. Recommended Values for VERIFY\_SUSPECT

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
timeout	2,000	5,000

**Table A.22. Recommended Values for pbcaster.NAKACK2**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
use_mcast_xmit	true	false
xmit_interval	1,000	Same as default
xmit_table_num_rows	50	50
xmit_table_msgs_per_row	10,000	1,024
xmit_table_max_compaction_time	10,000	30,000
max_msg_batch_size	100	Same as default
resend_last_seqno	false	true

**Table A.23. Recommended Values for UNICAST3**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
xmit_interval	500	Same as default
xmit_table_num_rows	100	50
xmit_table_msgs_per_row	1,0000	1,024
xmit_table_max_compaction_time	600,000	30,000
max_msg_batch_size	500	100
conn_close_timeout	60,000	No recommended value
conn_expiry_timeout	120,000	0

**Table A.24. Recommended Values for pbcaster.STABLE**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
stability_delay	6,000	500
desired_avg_gossip	20,000	5,000
max_bytes	2,000,000	1,000,000

**Table A.25. Recommended Values for pbcast.GMS**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
print_local_addr	true	false
join_timeout	5,000	15,000
view_bundling	true	Same as default

**Table A.26. Recommended Values for UFC**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
max_credits	500,000	2,000,000
min_threshold	0.40	Same as default

**Table A.27. Recommended Values for MFC**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
max_credits	500,000	2,000,000
min_threshold	0.40	Same as default

**Table A.28. Recommended Values for FRAG2**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
frag_size	60,000	Same as default

**Table A.29. Recommended Values for ENCRYPT**

Parameter	JGroups Default Value	JBoss Data Grid Configured Values
asymAlgorithm	RSA	-
asymInit	512	-
asymProvider	Bouncy Castle Provider	-
changeKeysOnViewChange	false	-
symAlgorithm	AES	-
symInit	128	-
symProvider	Bouncy Castle Provider	-

**NOTE**

See the *Red Hat JBoss Data Grid Developer Guide's JGroups ENCRYPT* section for details.

**Recommended Values for SASL**

See the *Red Hat JBoss Data Grid Developer Guide's User Authentication over Hot Rod Using SASL* section for details.

**Recommended Values for RELAY2**

See [Chapter 30, Set Up Cross-Datacenter Replication](#) for details.

[Report a bug](#)

**A.4. JBOSS DATA GRID JGROUPS CONFIGURATION FILES**

The following configuration files are included with JBoss Data Grid and contain the recommended values for JGroups:

- `default-configs/default-jgroups-ec2.xml`
- `default-configs/default-jgroups-google.xml`
- `default-configs/default-jgroups-tcp.xml`
- `default-configs/default-jgroups-udp.xml`

[Report a bug](#)

## APPENDIX B. CONNECTING WITH JCONSOLE

### B.1. CONNECT TO JDG VIA JCONSOLE

JConsole is a JMX GUI that allows a user to connect to a JVM, either local or remote, to monitor the JVM, its MBeans, and execute operations.

#### Procedure B.1. Add Management User to JBoss Data Grid

Before being able to connect to a remote JBoss Data Grid instance a user will need to be created; to add a user execute the following steps on the remote instance.

1. Navigate to the **bin** directory

```
cd $JDG_HOME/bin
```

2. Execute the **add-user.sh** script.

```
./add-user.sh
```

3. Accept the default option of **ManagementUser** by pressing return.
4. Accept the default option of **ManagementRealm** by pressing return.
5. Enter the desired username. In this example **jmxadmin** will be used.
6. Enter and confirm the password.
7. Accept the default option of no groups by pressing return.
8. Confirm that the desired user will be added to the **ManagementRealm** by entering **yes**.
9. Enter **no** as this user will not be used for connections between processes.
10. The following image shows an example execution run.

```
What type of user do you wish to add?
a) Management User (mgmt-users.properties)
b) Application User (application-users.properties)
(a):

Enter the details of the new user to add.
Realm (ManagementRealm) :
Username : jmxadmin
Password requirements are listed below. To modify these restrictions edit the add-user.properties configuration file.
- The password must not be one of the following restricted values {root, admin, administrator}
- The password must contain at least 8 characters, 1 alphabetic character(s), 1 digit(s), 1 non-alphanumeric symbol(s)
- The password must be different from the username
Password :
Re-enter Password :
What groups do you want this user to belong to? (Please enter a comma separated list, or leave blank for none)[ ]:
About to add user 'jmxadmin' for realm 'ManagementRealm'
Is this correct yes/no? yes
Added user 'jmxadmin' to file '/opt/jboss/jboss-datagrid-6.5.0-server/standalone/configuration/mgmt-users.properties'
Added user 'jmxadmin' with groups to file '/opt/jboss/jboss-datagrid-6.5.0-server/standalone/configuration/mgmt-groups.properties'
Is this new user going to be used for one AS process to connect to another AS process?
e.g. for a slave host controller connecting to the master or for a Remoting connection for server to server EJB calls.
yes/no? no
```

Figure B.1. Execution of add-user.sh

#### Binding the Management Interface

By default JBoss Data Grid will start with the management interface binding to 127.0.0.1. In order to connect remotely this interface must be bound to an IP address that is visible by the network. Either of the following options will correct this:

- **Option 1: Runtime** - By adjusting the `jboss.bind.address.management` property on startup a new IP address can be specified. In the following example JBoss Data Grid is starting with this bound to 192.168.122.5:

```
./standalone.sh ... -Djboss.bind.address.management=192.168.122.5
```

- **Option 2: Configuration** - Adjust the `jboss.bind.address.management` in the configuration file. This is found in the `interfaces` subsystem. A snippet of the configuration file, with the IP adjusted to 192.168.122.5, is provided below:

```
<interfaces>
  <interface name="management">
    <inet-address
value="{jboss.bind.address.management:192.168.122.5}"/>
    </interface>
  [...]
</interface>
```

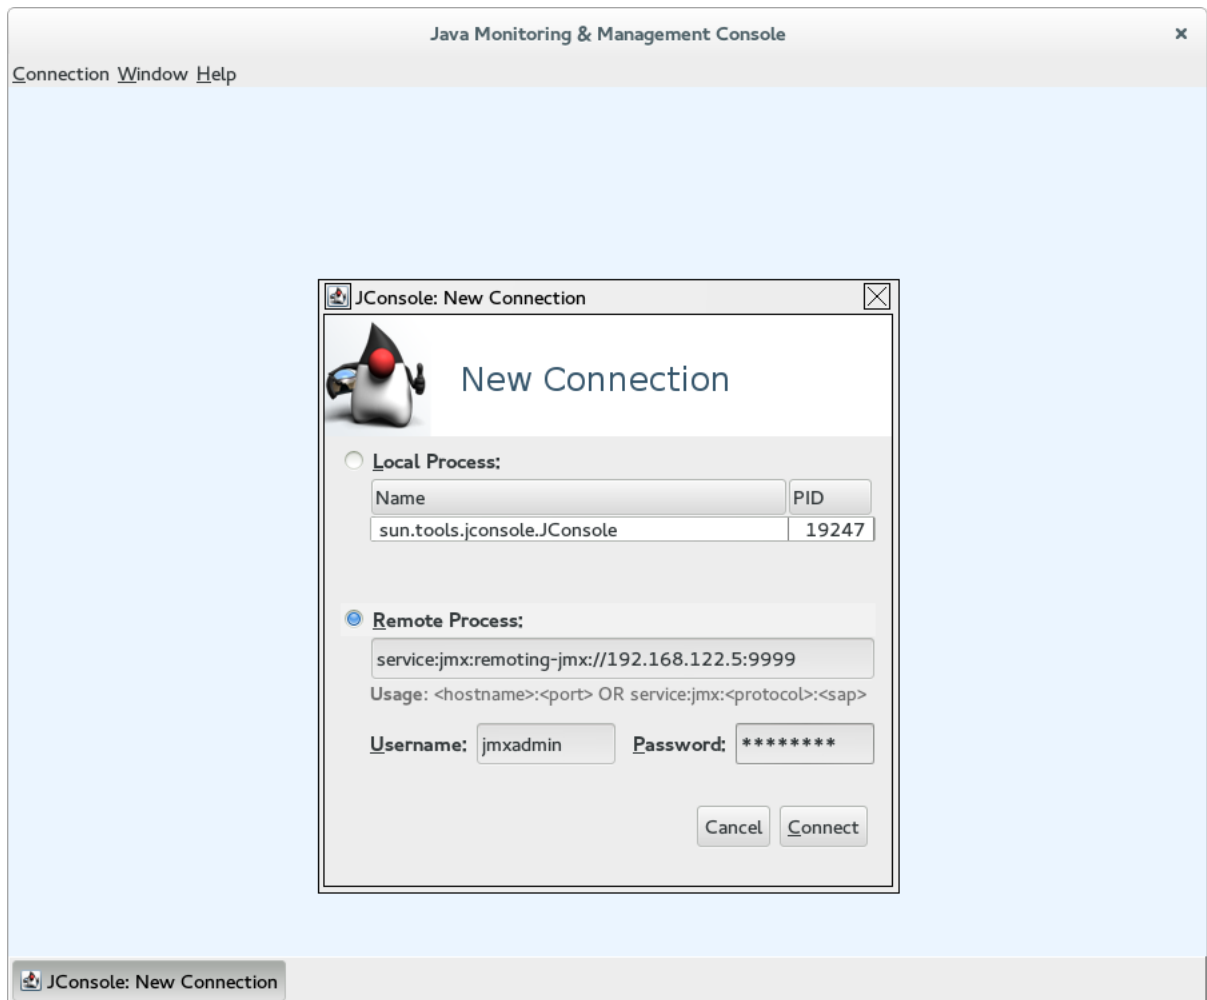
## Running JConsole

A `jconsole.sh` script is provided in the `$JDG_HOME/bin` directory. Executing this script will launch JConsole.

### Procedure B.2. Connecting to a remote JBoss Data Grid instance using JConsole

1. Execute the `$JDG_HOME/bin/jconsole.sh` script. This will result in the following window appearing:





**Figure B.2. JConsole**

2. Select **Remote Process**.
3. Enter **service:jmx:remoting-jmx://\$IP:9999** in the text area.
4. Enter the username and password, created from the **add-user.sh** script.
5. Click **Connect** to initiate the connection.
6. Once connected ensure that the cache-related nodes may be viewed. The following screenshot shows such a node.

Java Monitoring & Management Console

Connection Window Help

JBoss CLI / rhel6

Overview Memory Threads Classes VM Summary MBeans JBoss CLI / rhel6

+ JMImplementation  
 + com.sun.management  
 + java.lang  
 + java.nio  
 + java.util.logging  
 + jboss.as  
 + jboss.as.expr  
 - jboss.infinispan  
   - Cache  
     + "\_\_\_protobuf\_metadata(local)"  
     + "default(local)"  
     + "jboss-web-policy(local)"  
     + "memcachedCache(local)"  
     - "namedCache(local)"  
       - "local"  
         + Activation  
         + Cache  
         + ClusterCacheStats  
         + LockManager  
         + Passivation  
         + RollingUpgradeManager  
         - Statistics  
         + Attributes  
         + Operations  
         + "other(local)"  
     + CacheManager  
     + RemoteQuery

Attribute values

Name	Value
averageReadTime	0
averageRemoveTime	0
averageWriteTime	0
elapsedTime	92
evictions	0
hitRatio	0.0
hits	0
misses	0
numberOfEntries	0
readWriteRatio	0.0
removeHits	0
removeMisses	0
statisticsEnabled	true
stores	0
timeSinceReset	92

Refresh

JBoss CLI / rhel6

Figure B.3. JConsole: Showing a Cache

[Report a bug](#)

## APPENDIX C. JMX MBEANS IN REDHAT JBOSS DATA GRID

### C.1. ACTIVATION

`org.infinispan.eviction.ActivationManagerImpl`

Activates entries that have been passivated to the CacheStore by loading the entries into memory.

**Table C.1. Attributes**

Name	Description	Type	Writable
activations	Number of activation events.	String	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes

**Table C.2. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

### C.2. CACHE

`org.infinispan.CacheImpl`

The Cache component represents an individual cache instance.

**Table C.3. Attributes**

Name	Description	Type	Writable
cacheName	Returns the cache name.	String	No
cacheStatus	Returns the cache status.	String	No
configurationAsProperties	Returns the cache configuration in form of properties.	Properties	No

Name	Description	Type	Writable
version	Returns the version of Infinispan	String	No
cacheAvailability	Returns the cache availability	String	Yes

**Table C.4. Operations**

Name	Description	Signature
start	Starts the cache.	void start()
stop	Stops the cache.	void stop()
clear	Clears the cache.	void clear()

[Report a bug](#)

## C.3. CACHECONTAINERSTATS

`org.infinispan.stats.impl.CacheContainerStatsImpl`

The CacheContainerStats component contains statistics such as timings, hit/miss ratio, and operation information.

**Table C.5. Attributes**

Name	Description	Type	Writable
averageReadTime	Cache container total average number of milliseconds for all read operations in this cache container.	long	No
averageRemoveTime	Cache container total average number of milliseconds for all remove operations in this cache container.	long	No
averageWriteTime	Cache container total average number of milliseconds for all write operations in this cache container.	long	No

Name	Description	Type	Writable
evictions	Cache container total number of cache eviction operations.	long	No
hitRatio	Cache container total percentage hit/(hit+miss) ratio for this cache.	double	No
hits	Cache container total number of cache attribute hits.	long	No
misses	Cache container total number of cache attribute misses.	long	No
numberOfEntries	Cache container total number of entries currently in all caches from this cache container.	int	No
readWriteRatio	Cache container read/writes ratio in all caches from this cache container.	double	No
removeHits	Cache container total number of removal hits.	double	No
removeMisses	Cache container total number of cache removals where keys were not found.	long	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes
stores	Cache container total number of cache attribute put operations.	long	No

[Report a bug](#)

## C.4. CACHELOADER

`org.infinispan.interceptors.CacheLoaderInterceptor`

This component loads entries from a CacheStore into memory.

**Table C.6. Attributes**

Name	Description	Type	Writable
cacheLoaderLoads	Number of entries loaded from the cache store.	long	No
cacheLoaderMisses	Number of entries that did not exist in cache store.	long	No
stores	Returns a collection of cache loader types which are configured and enabled.	Collection	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes

**Table C.7. Operations**

Name	Description	Signature
disableStore	Disable all cache loaders of a given type, where type is a fully qualified class name of the cache loader to disable.	void disableStore(String storeType)
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.5. CACHEMANAGER

### `org.infinispan.manager.DefaultCacheManager`

The CacheManager component acts as a manager, factory, and container for caches in the system.

**Table C.8. Attributes**

Name	Description	Type	Writable
cacheManagerStatus	The status of the cache manager instance.	String	No

Name	Description	Type	Writable
clusterMembers	Lists members in the cluster.	String	No
clusterName	Cluster name.	String	No
clusterSize	Size of the cluster in the number of nodes.	int	No
createdCacheCount	The total number of created caches, including the default cache.	String	No
definedCacheCount	The total number of defined caches, excluding the default cache.	String	No
definedCacheNames	The defined cache names and their statuses. The default cache is not included in this representation.	String	No
name	The name of this cache manager.	String	No
nodeAddress	The network address associated with this instance.	String	No
physicalAddresses	The physical network addresses associated with this instance.	String	No
runningCacheCount	The total number of running caches, including the default cache.	String	No
version	Infinispan version.	String	No
globalConfigurationAsProperties	Global configuration properties	Properties	No

**Table C.9. Operations**

Name	Description	Signature
startCache	Starts the default cache associated with this cache manager.	void startCache()
startCache	Starts a named cache from this cache manager.	void startCache (String p0)

[Report a bug](#)

## C.6. CACHESTORE

### `org.infinispan.interceptors.CacheWriterInterceptor`

The CacheStore component stores entries to a CacheStore from memory.

**Table C.10. Attributes**

Name	Description	Type	Writable
writesToTheStores	Number of writes to the store.	long	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes

**Table C.11. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.7. CLUSTERCACHESTATS

### `org.infinispan.stats.impl.ClusterCacheStatsImpl`

The ClusterCacheStats component contains statistics such as timings, hit/miss ratio, and operation information for the whole cluster.

**Table C.12. Attributes**



Name	Description	Type	Writable
activations	The total number of activations in the cluster.	long	No
averageReadTime	Cluster wide total average number of milliseconds for a read operation on the cache.	long	No
averageRemoveTime	Cluster wide total average number of milliseconds for a remove operation in the cache.	long	No
averageWriteTime	Cluster wide average number of milliseconds for a write operation in the cache.	long	No
cacheLoaderLoads	The total number of cacheloader load operations in the cluster.	long	No
cacheLoaderMisses	The total number of cacheloader load misses in the cluster.	long	No
evictions	Cluster wide total number of cache eviction operations.	long	No
hitRatio	Cluster wide total percentage hit/(hit+miss) ratio for this cache.	double	No
hits	Cluster wide total number of cache hits.	long	No
invalidations	The total number of invalidations in the cluster.	long	No
misses	Cluster wide total number of cache attribute misses.	long	No
numberOfEntries	Cluster wide total number of entries currently in the cache.	int	No

Name	Description	Type	Writable
numberOfLocksAvailable	Total number of exclusive locks available in the cluster.	int	No
numberOfLocksHeld	The total number of locks held in the cluster.	int	No
passivations	The total number of passivations in the cluster.	long	No
readWriteRatio	Cluster wide read/writes ratio for the cache.	double	No
removeHits	Cluster wide total number of cache removal hits.	double	No
removeMisses	Cluster wide total number of cache removals where keys were not found.	long	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes
storeWrites	The total number of cachestore store operations in the cluster.	long	No
stores	Cluster wide total number of cache attribute put operations.	long	No
timeSinceStart	Number of seconds since the first cache node started.	long	No

**Table C.13. Operations**

Name	Description	Signature
setStaleStatsTreshold	Sets the threshold for cluster wide stats refresh (in milliseconds).	void setStaleStatsTreshold(long staleStatsThreshold)

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.8. DEADLOCKDETECTINGLOCKMANAGER

`org.infinispan.util.concurrent.locks.DeadlockDetectingLockManager`

This component provides information about the number of deadlocks that were detected.

**Table C.14. Attributes**

Name	Description	Type	Writable
detectedLocalDeadlocks	Number of local transactions that were rolled back due to deadlocks.	long	No
detectedRemoteDeadlocks	Number of remote transactions that were rolled back due to deadlocks.	long	No
overlapWithNotDeadlockAwareLockOwners	Number of situations when we try to determine a deadlock and the other lock owner is NOT a transaction. In this scenario we cannot run the deadlock detection mechanism.	long	No
totalNumberOfDetectedDeadlocks	Total number of local detected deadlocks.	long	No
concurrencyLevel	The concurrency level that the MVCC Lock Manager has been configured with.	int	No
numberOfLocksAvailable	The number of exclusive locks that are available.	int	No
numberOfLocksHeld	The number of exclusive locks that are held.	int	No

**Table C.15. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.9. DISTRIBUTIONMANAGER

### `org.infinispan.distribution.DistributionManagerImpl`

The `DistributionManager` component handles the distribution of content across a cluster.



#### NOTE

The `DistributionManager` component is only available in clustered mode.

**Table C.16. Operations**

Name	Description	Signature
isAffectedByRehash	Determines whether a given key is affected by an ongoing rehash.	boolean isAffectedByRehash(Object p0)
isLocatedLocally	Indicates whether a given key is local to this instance of the cache. Only works with String keys.	boolean isLocatedLocally(String p0)
locateKey	Locates an object in a cluster. Only works with String keys.	List locateKey(String p0)

[Report a bug](#)

## C.10. INTERPRETER

### `org.infinispan.cli.interpreter.Interpreter`

The `Interpreter` component executes command line interface (CLI operations).

**Table C.17. Attributes**

Name	Description	Type	Writable
cacheNames	Retrieves a list of caches for the cache manager.	String[]	No

**Table C.18. Operations**

Name	Description	Signature
createSessionId	Creates a new interpreter session.	String createSessionId(String cacheName)
execute	Parses and executes IspnCliQL statements.	String execute(String p0, String p1)

[Report a bug](#)

## C.11. INVALIDATION

### `org.infinispan.interceptors.InvalidatorInterceptor`

The Invalidation component invalidates entries on remote caches when entries are written locally.

**Table C.19. Attributes**

Name	Description	Type	Writable
invalidations	Number of invalidations.	long	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes

**Table C.20. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.12. LOCKMANAGER

### `org.infinispan.util.concurrent.locks.LockManagerImpl`

The LockManager component handles MVCC locks for entries.

**Table C.21. Attributes**

Name	Description	Type	Writable
------	-------------	------	----------

Name	Description	Type	Writable
concurrencyLevel	The concurrency level that the MVCC Lock Manager has been configured with.	int	No
numberOfLocksAvailable	The number of exclusive locks that are available.	int	No
numberOfLocksHeld	The number of exclusive locks that are held.	int	No

[Report a bug](#)

## C.13. LOCALTOPOLOGYMANAGER

### `org.infinispan.topology.LocalTopologyManagerImpl`

The `LocalTopologyManager` component controls the cache membership and state transfer in Red Hat JBoss Data Grid.



#### NOTE

The `LocalTopologyManager` component is only available in clustered mode.

**Table C.22. Attributes**

Name	Description	Type	Writable
rebalancingEnabled	If false, newly started nodes will not join the existing cluster nor will the state be transferred to them. If any of the current cluster members are stopped when rebalancing is disabled, the nodes will leave the cluster but the state will not be rebalanced among the remaining nodes. This will result in fewer copies than specified by the <b>numOwners</b> attribute until rebalancing is enabled again.	boolean	Yes

Name	Description	Type	Writable
clusterAvailability	If <b>AVAILABLE</b> the node is currently operating regularly. If <b>DEGRADED</b> then data can not be safely accessed due to either a network split, or successive nodes leaving.	String	No

[Report a bug](#)

## C.14. MASSINDEXER

`org.infinispan.query.MassIndexer`

The MassIndexer component rebuilds the index using cached data.

**Table C.23. Operations**

Name	Description	Signature
start	Starts rebuilding the index.	void start()



### NOTE

This operation is available only for caches with indexing enabled. For more information, see the Red Hat JBoss Data Grid *Infinispan Query Guide*

[Report a bug](#)

## C.15. PASSIVATION

`org.infinispan.eviction.PassivationManager`

The Passivation component handles the passivation of entries to a CacheStore on eviction.

**Table C.24. Attributes**

Name	Description	Type	Writable
passivations	Number of passivation events.	String	No
statisticsEnabled	Enables or disables the gathering of statistics by this component	boolean	Yes

**Table C.25. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.16. RECOVERYADMIN

`org.infinispan.transaction.xa.recovery.RecoveryAdminOperations`

The RecoveryAdmin component exposes tooling for handling transaction recovery.

**Table C.26. Operations**

Name	Description	Signature
forceCommit	Forces the commit of an in-doubt transaction.	String forceCommit(long p0)
forceCommit	Forces the commit of an in-doubt transaction	String forceCommit(int p0, byte[] p1, byte[] p2)
forceRollback	Forces the rollback of an in-doubt transaction.	String forceRollback(long p0)
forceRollback	Forces the rollback of an in-doubt transaction	String forceRollback(int p0, byte[] p1, byte[] p2)
forget	Removes recovery info for the given transaction.	String forget(long p0)
forget	Removes recovery info for the given transaction.	String forget(int p0, byte[] p1, byte[] p2)
showInDoubtTransactions	Shows all the prepared transactions for which the originating node crashed.	String showInDoubtTransactions()

[Report a bug](#)

## C.17. ROLLINGUPGRADEMANAGER

`org.infinispan.upgrade.RollingUpgradeManager`

The RollingUpgradeManager component handles the control hooks in order to migrate data from one version of Red Hat JBoss Data Grid to another.



**Table C.27. Operations**

Name	Description	Signature
disconnectSource	Disconnects the target cluster from the source cluster according to the specified migrator.	void disconnectSource(String p0)
recordKnownGlobalKeyset	Dumps the global known keyset to a well-known key for retrieval by the upgrade process.	void recordKnownGlobalKeyset()
synchronizeData	Synchronizes data from the old cluster to this using the specified migrator.	long synchronizeData(String p0)

[Report a bug](#)

## C.18. RPCMANAGER

`org.infinispan.remoting.rpc.RpcManagerImpl`

The `RpcManager` component manages all remote calls to remote cache instances in the cluster.



### NOTE

The `RpcManager` component is only available in clustered mode.

**Table C.28. Attributes**

Name	Description	Type	Writable
averageReplicationTime	The average time spent in the transport layer, in milliseconds.	long	No
committedViewAsString	Retrieves the committed view.	String	No
pendingViewAsString	Retrieves the pending view.	String	No
replicationCount	Number of successful replications.	long	No
replicationFailures	Number of failed replications.	long	No

Name	Description	Type	Writable
successRatio	Successful replications as a ratio of total replications.	String	No
successRatioFloatingPoint	Successful replications as a ratio of total replications in numeric double format.	double	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes

Table C.29. Operations

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()
setStatisticsEnabled	Whether statistics should be enabled or disabled (true/false)	void setStatisticsEnabled(boolean enabled)

[Report a bug](#)

## C.19. STATETRANSFERMANAGER

`org.infinispan.statetransfer.StateTransferManager`

The `StateTransferManager` component handles state transfer in Red Hat JBoss Data Grid.



### NOTE

The `StateTransferManager` component is only available in clustered mode.

Table C.30. Attributes

Name	Description	Type	Writable
joinComplete	If true, the node has successfully joined the grid and is considered to hold state. If false, the join process is still in progress..	boolean	No

Name	Description	Type	Writable
stateTransferInProgress	Checks whether there is a pending inbound state transfer on this cluster member.	boolean	No

[Report a bug](#)

## C.20. STATISTICS

### `org.infinispan.interceptors.CacheMgmtInterceptor`

This component handles general statistics such as timings, hit/miss ratio, etc.

**Table C.31. Attributes**

Name	Description	Type	Writable
averageReadTime	Average number of milliseconds for a read operation on the cache.	long	No
averageWriteTime	Average number of milliseconds for a write operation in the cache.	long	No
elapsedTime	Number of seconds since cache started.	long	No
evictions	Number of cache eviction operations.	long	No
hitRatio	Percentage hit/(hit+miss) ratio for the cache.	double	No
hits	Number of cache attribute hits.	long	No
misses	Number of cache attribute misses.	long	No
numberOfEntries	Number of entries currently in the cache.	int	No
readWriteRatio	Read/writes ratio for the cache.	double	No

Name	Description	Type	Writable
removeHits	Number of cache removal hits.	long	No
removeMisses	Number of cache removals where keys were not found.	long	No
stores	Number of cache attribute PUT operations.	long	No
timeSinceReset	Number of seconds since the cache statistics were last reset.	long	No
averageRemoveTime	Average number of milliseconds for a remove operation in the cache	long	No

**Table C.32. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.21. TRANSACTIONS

### `org.infinispan.interceptors.TxInterceptor`

The Transactions component manages the cache's participation in JTA transactions.

**Table C.33. Attributes**

Name	Description	Type	Writable
commits	Number of transaction commits performed since last reset.	long	No
prepares	Number of transaction prepares performed since last reset.	long	No

Name	Description	Type	Writable
rollbacks	Number of transaction rollbacks performed since last reset.	long	No
statisticsEnabled	Enables or disables the gathering of statistics by this component.	boolean	Yes

**Table C.34. Operations**

Name	Description	Signature
resetStatistics	Resets statistics gathered by this component.	void resetStatistics()

[Report a bug](#)

## C.22. TRANSPORT

### `org.infinispan.server.core.transport.NettyTransport`

The Transport component manages read and write operations to and from the server.

**Table C.35. Attributes**

Name	Description	Type	Writable
hostName	Returns the host to which the transport binds.	String	No
idleTimeout	Returns the idle timeout.	String	No
numberOfGlobalConnections	Returns a count of active connections in the cluster. This operation will make remote calls to aggregate results, so latency may have an impact on the speed of calculation for this attribute.	Integer	false
numberOfLocalConnections	Returns a count of active connections this server.	Integer	No

Name	Description	Type	Writable
numberWorkerThreads	Returns the number of worker threads.	String	No
port	Returns the port to which the transport binds.	String	
receiveBufferSize	Returns the receive buffer size.	String	No
sendBufferSize	Returns the send buffer size.	String	No
totalBytesRead	Returns the total number of bytes read by the server from clients, including both protocol and user information.	String	No
totalBytesWritten	Returns the total number of bytes written by the server back to clients, including both protocol and user information.	String	No
tcpNoDelay	Returns whether TCP no delay was configured or not.	String	No

[Report a bug](#)

## C.23. XSITEADMIN

### `org.infinispan.xsite.XSiteAdminOperations`

The XSiteAdmin component exposes tooling for backing up data to remote sites.

**Table C.36. Operations**

Name	Description	Signature
bringSiteOnline	Brings the given site back online on all the cluster.	String bringSiteOnline(String p0)
amendTakeOffline	Amends the values for 'TakeOffline' functionality on all the nodes in the cluster.	String amendTakeOffline(String p0, int p1, long p2)

Name	Description	Signature
getTakeOfflineAfterFailures	Returns the value of the 'afterFailures' for the 'TakeOffline' functionality.	String getTakeOfflineAfterFailures(String p0)
getTakeOfflineMinTimeToWait	Returns the value of the 'minTimeToWait' for the 'TakeOffline' functionality.	String getTakeOfflineMinTimeToWait(String p0)
setTakeOfflineAfterFailures	Amends the values for 'afterFailures' for the 'TakeOffline' functionality on all the nodes in the cluster.	String setTakeOfflineAfterFailures(String p0, int p1)
setTakeOfflineMinTimeToWait	Amends the values for 'minTimeToWait' for the 'TakeOffline' functionality on all the nodes in the cluster.	String setTakeOfflineMinTimeToWait(String p0, long p1)
siteStatus	Check whether the given backup site is offline or not.	String siteStatus(String p0)
status	Returns the status(offline/online) of all the configured backup sites.	String status()
takeSiteOffline	Takes this site offline in all nodes in the cluster.	String takeSiteOffline(String p0)
pushState	Starts the cross-site state transfer to the site name specified.	String pushState(String p0)
cancelPushState	Cancel the cross-site state transfer to the site name specified.	String cancelPushState(String p0)
getSendingSiteName	Returns the site name that is pushing state to this site.	String getSendingSiteName()
cancelReceiveState	Restores the site to the normal state. It is used when the link between the sites is broken during the state transfer.	String cancelReceiveState(String p0)
getPushStateStatus	Returns the status of completed and running cross-site state transfer.	String getPushStateStatus()
clearPushStateStatus	Clears the status of completed cross-site state transfer.	String clearPushStateStatus()

[Report a bug](#)



## APPENDIX D. CONFIGURATION RECOMMENDATIONS

### D.1. TIMEOUT VALUES

Table D.1. Timeout Value Recommendations for JBoss Data Grid

Timeout Value	Parent Element	Default Value	Recommended Value
distributedSyncTimeout	transport	240,000 (4 minutes)	Same as default
lockAcquisitionTimeout	locking	10,000 (10 seconds)	Same as default
cacheStopTimeout	transaction	30,000 (30 seconds)	Same as default
completedTxTimeout	transaction	60,000 (60 seconds)	Same as default
replTimeout	sync	15,000 (15 seconds)	Same as default
timeout	stateTransfer	240,000 (4 minutes)	Same as default
timeout	backup	10,000 (10 seconds)	Same as default
flushLockTimeout	async	1 (1 millisecond)	Same as default. Note that this value applies to asynchronous cache stores, but not asynchronous caches.
shutdownTimeout	async	25,000 (25 seconds)	Same as default. Note that this value applies to asynchronous cache stores, but not asynchronous caches.
pushStateTimeout	singletonStore	10,000 (10 seconds)	Same as default.
backup	replicationTimeout	10,000 (10 seconds)	
remoteCallTimeout	clusterLoader	0	For most requirements, same as default. This value is usually set to the same as the <b><i>sync.replTimeout</i></b> value.

[Report a bug](#)

## APPENDIX E. PERFORMANCE RECOMMENDATIONS

### E.1. CONCURRENT STARTUP FOR LARGE CLUSTERS

When starting a large number of instances, each managing a large number of caches, in parallel this may take a while as rebalancing attempts to distribute the data evenly as each node joins the cluster. To limit the number of rebalancing attempts made during the initial startup of the cluster disable rebalancing temporarily by following the below steps:

1. Start the first node in the cluster.
2. Set JMX attribute `jboss.infinispan/CacheManager/"clustered"/LocalTopologyManager/rebalancingEnabled` to `false`, as seen in [Section C.13](#), “LocalTopologyManager”.
3. Start the remaining nodes in the cluster.
4. Re-enable the JMX attribute `jboss.infinispan/CacheManager/"clustered"/LocalTopologyManager/rebalancingEnabled` by setting this value back to `true`, as seen in [Section C.13](#), “LocalTopologyManager”.

[Report a bug](#)

## APPENDIX F. REFERENCES

### F.1. ABOUT CONSISTENCY

Consistency is the policy that states whether it is possible for a data record on one node to vary from the same data record on another node.

For example, due to network speeds, it is possible that a write operation performed on the master node has not yet been performed on another node in the store, a strong consistency guarantee will ensure that data which is not yet fully replicated is not returned to the application.

[Report a bug](#)

### F.2. ABOUT CONSISTENCY GUARANTEE

Despite the locking of a single owner instead of all owners, Red Hat JBoss Data Grid's consistency guarantee remains intact. Consider the following situation:

1. If Key **K** is hashed to nodes **{A, B}** and transaction **TX1** acquires a lock for **K** on, for example, node **A** and
2. If another cache access occurs on node **B**, or any other node, and **TX2** attempts to lock **K**, this access attempt fails with a timeout because the transaction **TX1** already holds a lock on **K**.

This lock acquisition attempt always fails because the lock for key **K** is always deterministically acquired on the same node of the cluster, irrespective of the transaction's origin.

[Report a bug](#)

### F.3. ABOUT JBOSS CACHE

Red Hat JBoss Cache is a tree-structured, clustered, transactional cache that can also be used in a standalone, non-clustered environment. It caches frequently accessed data in-memory to prevent data retrieval or calculation bottlenecks that occur while enterprise features such as Java Transactional API (JTA) compatibility, eviction and persistence are provided.

JBoss Cache is the predecessor to Infinispan and Red Hat JBoss Data Grid.

[Report a bug](#)

### F.4. ABOUT RELAY2

The **RELAY** protocol bridges two remote clusters by creating a connection between one node in each site. This allows multicast messages sent out in one site to be relayed to the other and vice versa.

JGroups includes the **RELAY2** protocol, which is used for communication between sites in Red Hat JBoss Data Grid's Cross-Site Replication.

The **RELAY2** protocol works similarly to **RELAY** but with slight differences. Unlike **RELAY**, the **RELAY2** protocol:

- connects more than two sites.
- connects sites that operate autonomously and are unaware of each other.

- offers both unicasts and multicast routing between sites.

[Report a bug](#)

## F.5. ABOUT RETURN VALUES

Values returned by cache operations are referred to as return values. In Red Hat JBoss Data Grid, these return values remain reliable irrespective of which cache mode is employed and whether synchronous or asynchronous communication is used.

[Report a bug](#)

## F.6. ABOUT RUNNABLE INTERFACES

A Runnable Interface (also known as a Runnable) declares a single `run()` method, which executes the active part of the class' code. The Runnable object can be executed in its own thread after it is passed to a thread constructor.

[Report a bug](#)

## F.7. ABOUT TWO PHASE COMMIT (2PC)

A Two Phase Commit protocol (2PC) is a consensus protocol used to atomically commit or roll back distributed transactions. It is successful when faced with cases of temporary system failures, including network node and communication failures, and is therefore widely utilized.

[Report a bug](#)

## F.8. ABOUT KEY-VALUE PAIRS

A key-value pair (KVP) is a set of data consisting of a key and a value.

- A key is unique to a particular data entry. It consists of entry data attributes from the related entry.
- A value is the data assigned to and identified by the key.

[Report a bug](#)

## F.9. THE EXTERNALIZER

### F.9.1. About Externalizer

An `Externalizer` is a class that can:

- Marshall a given object type to a byte array.
- Unmarshall the contents of a byte array into an instance of the object type.

Externalizers are used by Red Hat JBoss Data Grid and allow users to specify how their object types are serialized. The marshalling infrastructure used in JBoss Data Grid builds upon JBoss Marshalling and provides efficient payload delivery and allows the stream to be cached. The stream caching allows data to be accessed multiple times, whereas normally a stream can only be read once.

[Report a bug](#)

## F.9.2. Internal Externalizer Implementation Access

Externalizable objects should not access Red Hat JBoss Data Grids Externalizer implementations. The following is an example of incorrect usage:

```
public static class ABCMarshallingExternalizer implements
AdvancedExternalizer<ABCMarshalling> {
    @Override
    public void writeObject(ObjectOutput output, ABCMarshalling object)
throws IOException {
        MapExternalizer ma = new MapExternalizer();
        ma.writeObject(output, object.getMap());
    }

    @Override
    public ABCMarshalling readObject(ObjectInput input) throws IOException,
ClassNotFoundException {
        ABCMarshalling hi = new ABCMarshalling();
        MapExternalizer ma = new MapExternalizer();
        hi.setMap((ConcurrentHashMap<Long, Long>) ma.readObject(input));
        return hi;
    }
    <!-- Additional configuration information here -->
}
```

End user externalizers do not need to interact with internal externalizer classes. The following is an example of correct usage:

```
public static class ABCMarshallingExternalizer implements
AdvancedExternalizer<ABCMarshalling> {
    @Override
    public void writeObject(ObjectOutput output, ABCMarshalling object)
throws IOException {
        output.writeObject(object.getMap());
    }

    @Override
    public ABCMarshalling readObject(ObjectInput input) throws IOException,
ClassNotFoundException {
        ABCMarshalling hi = new ABCMarshalling();
        hi.setMap((ConcurrentHashMap<Long, Long>) input.readObject());
        return hi;
    }
    <!-- Additional configuration information here -->
}
```

[Report a bug](#)

## F.10. HASH SPACE ALLOCATION

### F.10.1. About Hash Space Allocation

Red Hat JBoss Data Grid is responsible for allocating a portion of the total available hash space to each node. During subsequent operations that must store an entry, JBoss Data Grid creates a hash of the relevant key and stores the entry on the node that owns that portion of hash space.

[Report a bug](#)

## F.10.2. Locating a Key in the Hash Space

Red Hat JBoss Data Grid always uses an algorithm to locate a key in the hash space. As a result, the node that stores the key is never manually specified. This scheme allows any node to know which node owns a particular key without such ownership information being distributed. This scheme reduces the amount of overhead and, more importantly, improves redundancy because the ownership information does not need to be replicated in case of node failure.

[Report a bug](#)

## F.10.3. Requesting a Full Byte Array

### How can I request the Red Hat JBoss Data Grid return a full byte array instead of partial byte array contents?

As a default, JBoss Data Grid only partially prints byte arrays to logs to avoid unnecessarily printing large byte arrays. This occurs when either:

- JBoss Data Grid caches are configured for lazy deserialization. Lazy deserialization is not available in JBoss Data Grid's Remote Client-Server mode.
- A **Memcached** or **Hot Rod** server is run.

In such cases, only the first ten positions of the byte array display in the logs. To display the complete contents of the byte array in the logs, pass the ***-Dinfinispan.arrays.debug=true*** system property at start up.

### Example F.1. Partial Byte Array Log

```
2010-04-14 15:46:09,342 TRACE [ReadCommittedEntry] (HotRodWorker-1-1)
Updating entry
(key=CacheKey{data=ByteArray{size=19, hashCode=1b3278a,
array=[107, 45, 116, 101, 115, 116, 82, 101, 112, 108, ..]}}
removed=false valid=true changed=true created=true
value=CacheValue{data=ByteArray{size=19,
array=[118, 45, 116, 101, 115, 116, 82, 101, 112, 108, ..]},
version=281483566645249}]
And here's a log message where the full byte array is shown:
2010-04-14 15:45:00,723 TRACE [ReadCommittedEntry] (Incoming-
2,Infinispan-Cluster,eq-6834) Updating entry
(key=CacheKey{data=ByteArray{size=19, hashCode=6cc2a4,
array=[107, 45, 116, 101, 115, 116, 82, 101, 112, 108, 105, 99, 97, 116,
101, 100, 80, 117, 116]}}
removed=false valid=true changed=true created=true
value=CacheValue{data=ByteArray{size=19,
array=[118, 45, 116, 101, 115, 116, 82, 101, 112, 108, 105, 99, 97, 116,
101, 100, 80, 117, 116]},
version=281483566645249}]
```

[Report a bug](#)

## APPENDIX G. REVISION HISTORY

<b>Revision 6.5.0-12</b>	<b>Thu Mar 10 2016</b>	<b>Christian Huffman</b>
Included section on Google Compute Engine.		
<b>Revision 6.5.0-11</b>	<b>Wed Feb 24 2016</b>	<b>Christian Huffman</b>
BZ-1310604: Clarified the versions in which externalization of HTTP sessions is supported.		
<b>Revision 6.5.0-10</b>	<b>Mon Oct 26 2015</b>	<b>Rakesh Ghatvisave, Christian Huffman</b>
BZ-1262366: Added Hot Rod clients and server compatibility table.		
BZ-1179998: Included information on detecting and recovering from a split brain timeout.		
BZ-1267168: Updated verbage.		
BZ-1265297: Added clarification on state transfer.		
<b>Revision 6.5.0-9</b>	<b>Mon Sep 14 2015</b>	<b>Christian Huffman</b>
Updating for 6.5.1 release.		
<b>Revision 6.5.0-8</b>	<b>Wed Jun 24 2015</b>	<b>Christian Huffman</b>
BZ-1191416: Updated JPA Cache Store information.		
<b>Revision 6.5.0-7</b>	<b>Wed Jun 3 2015</b>	<b>Christian Huffman</b>
BZ-1188924: Added information on starting a large number of caches in parallel.		
BZ-1188922: Included JMX metrics for ClusterCacheStats and CacheContainerStats.		
BZ-1188906: Updated Hot Rod C++ client section.		
BZ-1227899: Added section on externalizing HTTP sessions from EAP.		
BZ-1228226: Added instructions on generating a maven archetype for custom cache stores.		
<b>Revision 6.5.0-6</b>	<b>Wed May 27 2015</b>	<b>Rakesh Ghatvisave</b>
BZ-1203175: Updated state transfer between sites topic.		
<b>Revision 6.5.0-5</b>	<b>Mon May 18 2015</b>	<b>Rakesh Ghatvisave</b>
BZ-1188906: Updated Hot Rod C++ client support platform.		
<b>Revision 6.5.0-4</b>	<b>Fri May 08 2015</b>	<b>Christian Huffman</b>
BZ-1200763: Added clusterAvailability in LocalTopologyManager MBean.		
<b>Revision 6.5.0-3</b>	<b>Wed 06 May 2015</b>	<b>Christian Huffman</b>
Restructured adding a Custom Cache Store.		
<b>Revision 6.5.0-2</b>	<b>Mon 04 May 2015</b>	<b>Rakesh Ghatvisave</b>
BZ-1188928: Added creating cache using JON topic.		
BZ-1188921: Added C# string marshaller topic.		
<b>Revision 6.5.0-1</b>	<b>Mon 27 Apr 2015</b>	<b>Rakesh Ghatvisave</b>
BZ-1188905: Added near caching topic.		
<b>Revision 6.5.0-0</b>	<b>Mon 13 Apr 2015</b>	<b>Rakesh Ghatvisave</b>
BZ-1188921: Removed admonition notes about C# Tech Preview mode.		