



# Red Hat JBoss Data Grid 7.2

## 7.2.2 Release Notes

Release Information for JBoss Data Grid



# Red Hat JBoss Data Grid 7.2 7.2.2 Release Notes

---

Release Information for JBoss Data Grid

## Legal Notice

Copyright © 2018 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

JBoss Data Grid 7.2.2 Release Notes describe new features and enhancements. These Release Notes also provide details about issues that exist in this release of JBoss Data Grid as well as issues that this release resolves, along with information on supported configurations and component versions.

---

## Table of Contents

<b>CHAPTER 1. JBOSS DATA GRID 7.2.2</b> .....	<b>3</b>
1.1. INTRODUCTION TO RED HAT JBOSS DATA GRID	3
1.2. ABOUT THE RELEASE NOTES	3
<b>CHAPTER 2. PATCHING RED HAT JBOSS DATA GRID SERVER INSTANCES</b> .....	<b>4</b>
2.1. APPLYING PATCHES TO RED HAT JBOSS DATA GRID SERVER INSTANCES	4
2.2. REVERTING PATCHES	4
<b>CHAPTER 3. NEW FEATURES AND ENHANCEMENTS</b> .....	<b>6</b>
3.1. JMX MBEANS PROVIDE STATISTICS IN NANOSECONDS	6
3.2. STATISTICS IMPROVEMENTS	6
3.3. CONFIGURABLE TIMEOUTS FOR OPERATIONS WITH SPRING CACHE	6
3.4. HOT ROD NODE.JS CLIENT ENHANCEMENTS	6
3.5. TCP KEEPALIVE CONFIGURATION PARAMETER FOR SERVER CONNECTIONS	6
3.6. PERFORMANCE IMPROVEMENTS TO GENERATED ICKLE QUERIES	7
<b>CHAPTER 4. SUPPORTED CONFIGURATIONS</b> .....	<b>8</b>
4.1. SUPPORTED CONFIGURATIONS	8
<b>CHAPTER 5. COMPONENT VERSIONS</b> .....	<b>9</b>
5.1. COMPONENT VERSIONS	9
<b>CHAPTER 6. KNOWN ISSUES</b> .....	<b>10</b>
6.1. KNOWN ISSUES	10
<b>CHAPTER 7. RESOLVED ISSUES</b> .....	<b>13</b>
7.1. RESOLVED ISSUES	13



## CHAPTER 1. JBOSS DATA GRID 7.2.2



### IMPORTANT

A known issue exists with the patch for JBoss Data Grid 7.2.2. You should not apply this patch to your JBoss Data Grid deployment. Use the patch for JBoss Data Grid 7.2.3 or later instead. Patches are cumulative so any issues or features in JBoss Data Grid 7.2.2 are included in subsequent releases.

See [Known Issues](#) for more information.

### 1.1. INTRODUCTION TO RED HAT JBOSS DATA GRID

JBoss Data Grid is an open source, distributed, in-memory key/value data store built from the Infinispan open source software project. JBoss Data Grid is designed to provide elastic, high performance, highly available capabilities that scale linearly.

You can deploy JBoss Data Grid as a server in Remote Client-Server mode or embedded in a Java Virtual Machine (JVM) in Library mode to distribute and replicate data across a cluster of nodes.

JBoss Data Grid allows you to optionally write data to disk and provides access to data from a variety of clients as well as through a RESTful API, the Memcached and Hot Rod protocols, or directly in process through a Java Map API.

### 1.2. ABOUT THE RELEASE NOTES

Release Notes provide information about new features as well as known and resolved issues. You should read this document before installing JBoss Data Grid 7.2.

You can find the complete set of documentation for JBoss Data Grid at the [JBoss Data Grid Product Documentation](#) page in the Red Hat Customer Service Portal.

## CHAPTER 2. PATCHING RED HAT JBOSS DATA GRID SERVER INSTANCES

### 2.1. APPLYING PATCHES TO RED HAT JBOSS DATA GRID SERVER INSTANCES

Red Hat JBoss Data Grid server uses the patching functionality in JBoss Enterprise Application Platform (EAP) so that you can apply changes from errata releases without the need to completely replace an existing installation.

Patches are distributed for cumulative updates within a release version. The base version for this release is 7.2.0. You can apply 7.2.x patches to the base version or on top of other patches.

You cannot apply 7.2.x patches to any other JBoss Data Grid release version. Likewise you cannot apply patches from other release versions to the 7.2 release.

JBoss Data Grid provides patches for server instances only (Remote Client-Server mode). All other distributions, such as EAP modules, clients, and JBoss Data Grid Library mode, are provided as full releases.



#### IMPORTANT

A known issue exists with the patch for JBoss Data Grid 7.2.2. You should not apply this patch to your JBoss Data Grid deployment. Use the patch for JBoss Data Grid 7.2.3 or later instead. See [Known Issues](#) for more information.

To apply a patch to JBoss Data Grid, do the following:

1. Download the patch from the Red Hat Customer Portal at <https://access.redhat.com/downloads/>
2. Stop the server instance that you want to patch if it is running.  
To avoid issues with classloading, you should not apply patches to JBoss Data Grid while the server is running.

Either use the Administration Console to stop the server or enter **Ctrl-C** in the terminal where JBoss Data Grid is running.

3. Open a terminal and change to the **JDG\_HOME** directory.

```
$ cd JDG_HOME
```

4. Apply the patch as follows:

```
$ bin/cli.sh "patch apply /path/to/jboss-datagrid-7.2.x-server-patch.zip"
```

5. Start the server with either the **standalone.sh** or **domain.sh** script, for example:

```
$ bin/standalone.sh -c clustered.xml
```

### 2.2. REVERTING PATCHES

You can roll back patches to revert the Red Hat JBoss Data Grid server to the previously installed version.



### IMPORTANT

You should roll back patches only after applying a patch that results in unexpected behavior or undesirable effects. Rolling back patches is not intended for general uninstall functionality.

To revert a JBoss Data Grid patch, do the following:

1. Stop the server instance that you want to roll back if it is running. Either use the Administration Console to stop the server or enter **Ctrl-C** in the terminal where JBoss Data Grid is running.

2. Open a terminal and change to the **JDG\_HOME** directory.

```
$ cd JDG_HOME
```

3. Find the ID of the patch that you want to roll back.

```
$ bin/cli.sh "patch history"
```

4. Roll back the server version as follows:

```
$ bin/cli.sh "patch rollback --patch-id=PATCH_ID --reset-configuration=false"
```



### WARNING

Use caution when specifying the **reset-configuration** option.

**--reset-configuration=false** does not revert the server configuration. Because applying patches can change the server configuration, it is possible that the server does not restart if you roll back the patch but do not roll back the configuration. In this case, you should verify the server configuration and manually adjust it as needed before starting the server.

**--reset-configuration=true** reverts the server configuration to the pre-patch state. Any changes to the server configuration after the patch was applied are removed.

If conflicts exist when you attempt to roll back the patch, the operation fails and warnings occur. Enter **patch --help** to list available arguments that you can use to resolve the conflicts.

5. Start the server with either the **standalone.sh** or **domain.sh** script.

## CHAPTER 3. NEW FEATURES AND ENHANCEMENTS

### 3.1. JMX MBEANS PROVIDE STATISTICS IN NANOSECONDS

The following JMX MBeans in JBoss Data Grid now include statistics about read and write operations in nanoseconds:

- **CacheContainerStats**
- **ClusterCacheStats**
- **Statistics**

For more information, see [JMX MBeans](#) in the Administration and Configuration Guide.

### 3.2. STATISTICS IMPROVEMENTS

Several improvements have been made to statistics that are exposed via JMX MBeans in JBoss Data Grid.

- The **dataMemoryUsed** statistic shows the amount of memory, in bytes, for a cache or cache container when eviction is enabled. This statistic is exposed through the **CacheContainerStats** and **ClusterCacheStats** MBeans.
- The **requiredMinimumNumberOfNodes** statistic shows the number of nodes that are required to guarantee data consistency. This statistic is exposed through the **Statistics** MBean.

For more information, see [JMX MBeans](#) in the Administration and Configuration Guide.

### 3.3. CONFIGURABLE TIMEOUTS FOR OPERATIONS WITH SPRING CACHE

When using JBoss Data Grid as a Spring Cache provider, you can now configure timeouts for read and write operations on the cache with the following properties:

- **infinispan.spring.operation.read.timeout** specifies the time, in milliseconds, to wait for read operations to complete. The default is **0** which means unlimited wait time.
- **infinispan.spring.operation.write.timeout** specifies the time, in milliseconds, to wait for write operations to complete. The default is **0** which means unlimited wait time.

For more information, see [Integration with the Spring Framework](#) in the Developer Guide.

### 3.4. HOT ROD NODE.JS CLIENT ENHANCEMENTS

This release improves the Hot Rod Node.js client by including memory profiling capabilities and enhancements to avoid excessive memory consumption.

For more information, see [Memory Profiling with the Hot Rod Node.js Client](#) in the Developer Guide.

### 3.5. TCP KEEPALIVE CONFIGURATION PARAMETER FOR SERVER CONNECTIONS

This release includes a `tcpKeepAlive()` method in the `org.infinispan.server.core.configuration.ProtocolServerConfigurationChildBuilder` interface. You can use this method to configure JBoss Data Grid server to keep TCP connections active. The default is to disable TCP keepalive.

### 3.6. PERFORMANCE IMPROVEMENTS TO GENERATED ICKLE QUERIES

The `org.infinispan.query.dsl.QueryBuilder` interface has been improved so that the Ickle queries it generates are easier to parse.

## CHAPTER 4. SUPPORTED CONFIGURATIONS

### 4.1. SUPPORTED CONFIGURATIONS

Supported hardware and software configurations for JDG 7.2 are available on the Customer Portal at <https://access.redhat.com/articles/2435931>.

## CHAPTER 5. COMPONENT VERSIONS

### 5.1. COMPONENT VERSIONS

The full list of component versions used in Red Hat JBoss Data Grid 7.2 are available on the Customer Portal at <https://access.redhat.com/site/articles/488833>.

## CHAPTER 6. KNOWN ISSUES

### 6.1. KNOWN ISSUES

Find out about known issues and available workarounds in this release of Red Hat JBoss Data Grid.

#### Cannot Apply Patches to JBoss Data Grid Server If the 7.2.2 Patch Is Applied

**Issue:** [JDG-2178](#)

**Description:** If you apply the patch for JBoss Data Grid 7.2.2, you cannot apply patches for subsequent releases.

**Workaround:** Do not apply the patch for JBoss Data Grid 7.2.2. If you have already applied the patch and plan to apply a patch for a later release, you must first revert the patch for JBoss Data Grid 7.2.2 as follows:

```
$ bin/cli.sh "patch rollback --patch-id=jboss-jdg-7.2.2.CP --reset-configuration=false"
```

#### JBoss Data Grid Quickstarts Missing Dependencies

**Issues:** [JDG-2100](#) and [JDG-2097](#)

**Description:** The `secure-embedded-cache` and `cdi-jdg` quickstarts fail with errors when you attempt to run them.

This issue is the result of missing dependencies for the quickstarts.

**Workaround:** Update dependencies in the quickstarts before you run them.

To update the `secure-embedded-cache` quickstart, add `org.infinispan` to the `maven-war-plugin` artifact in your `pom.xml` file. You can find the correct dependency [on line 134 in the pom.xml on Github](#).

To update the `cdi-jdg` quickstart, add the `infinispan-jcache` artifact to your `pom.xml` file. You can find this dependency [on line 97 in the pom.xml on Github](#).

#### JBoss Data Grid Does Not Start and NullPointerException Occurs if Persistent Location is Not Unique

**Issue:** [JDG-1504](#)

**Description:** It is not possible to start more than one instance of JBoss Data Grid from the same directory if the value of `persistent-location` is not unique in each configuration.

**Workaround:** Do one of the following to ensure the value of `persistent-location` is unique in the JBoss Data Grid configuration:

- Set a unique `path` attribute for each instance.

```
<persistent-location path="mydir${uniquesuffix}" />
```

- Set a unique `relative-to` attribute for each instance if using a common path:

```
<paths>
  <path name="mypath" path="/home/user/jboss-datagrid-7.2.0-server/standalone/mypath" />
</paths>
```

```
</paths>
```

```
<persistent-location relative-to="mypath" path="mydir" />
```

- Specify the `-Djboss.server.data.dir` system property to set a unique relative location for each instance if using a common path:

```
./standalone.sh -Djboss.server.data.dir=/home/user/jboss-datagrid-7.2.0-server/standalone/mypath
```

### SKIP\_CACHE\_LOAD Flag Has No Effect if Authentication is Enabled

**Issue:** [JDG-1424](#)

**Description:** In Remote Client-Server mode, if you set the `SKIP_CACHE_LOAD` flag in the cache store configuration and enable authentication on Hot Rod clients, all entries are retrieved from the cache, including evicted entries.

**Workaround:** There is no workaround for this issue.

### Intermittent Data Loss Occurs During Rolling Upgrades Between Clusters

**Issue:** [JDG-991](#)

**Description:** When performing a rolling upgrade of JBoss Data Grid, all migrated data can be deleted from the target cluster after the nodes in the source cluster are taken offline.

**Workaround:** There is no workaround for this issue.

### Cluster Actions Disabled on JBoss Data Grid Administration Console in Reload-Required State

**Issue:** [JDG-1843](#)

**Description:** Actions available for the JBoss Data Grid cluster are not available in the Administration Console if you choose to restart the cluster after changing the configuration. In this case, the cluster is in the **Reload-Required** state.

**Reload** and **Stop** actions are available for each node in the cluster.

**Workaround:** Reload at least one node in the cluster to restore actions at the cluster level.

### Errors Occur When Changing the Eviction Strategy from the JBoss Data Grid Administration Console

**Issue:** [JDG-1804](#)

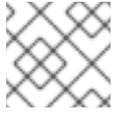
**Description:** If JBoss Data Grid is running in domain mode and you change the eviction strategy in the configuration through the Administration Console but do not restart to apply the changes, an error occurs.

**Workaround:** Restart the server after changing to the eviction strategy.

### NullPointerException Occurs When Reading Data from Persistent Storage in JBoss Data Grid 7.0 and Earlier

**Issue:** [JDG-968](#)

**Description:** If you store data in a cache store with JBoss Data Grid 7.0 and earlier and then attempt to read that data with JBoss Data Grid 7.1 or later, an error occurs and it is not possible to read the data.



**NOTE**

This issue does not apply when upgrading from JBoss Data Grid 7.1 to 7.2.

**Workaround:** There is no workaround for this issue.

## CHAPTER 7. RESOLVED ISSUES

### 7.1. RESOLVED ISSUES

This release of JBoss Data Grid resolves the following issues:

#### **JDG-2007 - Near-Real-Time Indexes Lost when Server is Stopped**

When using JBoss Data Grid in Remote Client-Server Mode, cache indexes that used Near-Real-Time indexing capabilities were not persisted to storage during shutdown. As a result, those cache indexes were lost.

#### **JDG-1835 - Disabling Indexing for Types with the `@Indexed(false)` Annotation Creates Empty Documents**

Protobuf encoded entries of message types that had the `@Indexed(false)` annotation resulted in empty documents being added to the cache index.

#### **JDG-2026 - Null Pointer Exceptions Occur with Nested Protobuf Messages**

JBoss Data Grid threw a null pointer exception when Protobuf encoded messages included the `@Indexed(false)` annotation but also contained nested messages without any annotation.

#### **JDG-1975 - Hot Rod Clients Should Consistently Wrap `InterruptedException` in `HotRodClientException`**

Hot Rod clients did not wrap the client-level `InterruptedException` exception in `HotRodClientException`.

#### **JDG-2027 - Hot Rod and REST Interoperability Does Not Work with Non-String Keys**

When requesting entries via the REST API from a cache that stored Protobuf encoded entries, the `Key-Content-Type` request header was ignored and an HTTP 404 error was returned.

#### **JDG-2043 - Lock Timeout when Registering a Queried Class to the Registry Cache**

JBoss Data Grid behaved unexpectedly in some cases where indexing was enabled and classes were added to an internal cache that was populated by multiple clients.

#### **JDG-2041 - Null Pointer Exception for Queries if Partition Handling Enabled and Cache in a Degraded State**

JBoss Data Grid returned a null pointer exception for queries to caches where partition handling was enabled and not configured to use the `ALLOW_READ_WRITES` strategy. In this case, JBoss Data Grid operates in a degraded state and indexes are not consistent.

#### **JDG-2044 - No Exception Occurs for Queries with Caches in a Degraded State**

JBoss Data Grid now returns the `AvailabilityException` exception for queries to caches where partition handling is enabled and not configured to use the `ALLOW_READ_WRITES` strategy.

#### **JDG-2050 - Caches Created via Hot Rod Ignore Media Type Configuration**

The media type configuration did not take effect for caches that were created via the Hot Rod interface.

#### **JDG-2052 - Deprecation Annotations for `DistributionManager`**

Deprecation annotations were incorrectly applied to the `org.infinispan.distribution.DistributionManager` interface.