



Red Hat JBoss Data Grid 7.2

7.2.1 Release Notes

Release Information for JBoss Data Grid

Red Hat JBoss Data Grid 7.2 7.2.1 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 7.2.1 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

CHAPTER 1. RED HAT JBOSS DATA GRID 7.2.1	3
1.1. INTRODUCTION TO RED HAT JBOSS DATA GRID	3
1.2. ABOUT THE RELEASE NOTES	3
CHAPTER 2. PATCHING RED HAT JBOSS DATA GRID SERVER INSTANCES	4
2.1. APPLYING PATCHES TO RED HAT JBOSS DATA GRID SERVER INSTANCES	4
2.2. REVERTING PATCHES	4
CHAPTER 3. NEW FEATURES AND ENHANCEMENTS	6
3.1. CONFIGURABLE WHITELIST FOR DESERIALIZATION	6
3.2. DISABLING INDEXING FOR PROTOBUF SCHEMA FILES	6
3.3. EXPOSE QUERY.GETQUERYSTRING() AS OFFICIAL API	6
CHAPTER 4. SUPPORTED CONFIGURATIONS	7
4.1. SUPPORTED CONFIGURATIONS	7
CHAPTER 5. COMPONENT VERSIONS	8
5.1. COMPONENT VERSIONS	8
CHAPTER 6. KNOWN ISSUES	9
6.1. KNOWN ISSUES	9
CHAPTER 7. RESOLVED ISSUES	11
7.1. RESOLVED ISSUES	11

CHAPTER 1. RED HAT JBOSS DATA GRID 7.2.1

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

To avoid issues with classloading, you should not apply patches to JBoss Data Grid while the server is running.

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.
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. CONFIGURABLE WHITELIST FOR DESERIALIZATION

By default the JBoss Data Grid server allows deserialization only for strings and primitives. As of this release, you can specify other Java class instances in a whitelist so that JBoss Data Grid deserializes objects that belong to those classes.



NOTE

If you do not add your application classes to the whitelist, JBoss Data Grid does not deserialize objects that belong to those classes.

On the client side, you can add a whitelist to restrict deserialization to objects that belong to specific Java classes.

For more information, see the following:

- [Configuring the Deserialization Whitelist](#) in the Administration and Configuration Guide describes how to configure the JBoss Data Grid server.
- [Restricting Deserialization to Specific Java Classes](#) in the Developer Guide describes how to configure clients.

3.2. DISABLING INDEXING FOR PROTOBUF SCHEMA FILES

After you enable indexing for a cache, all fields of Protobuf encoded entries are indexed by default. This indexing can be inefficient when handling Protobuf message types that have many or large fields. You can now disable indexing for all Protobuf message types that are not annotated with the `indexed_by_default` option. For more information, see [Disabling Indexing for All Protobuf Message Types](#) in the Developer Guide.

3.3. EXPOSE QUERY.GETQUERYSTRING() AS OFFICIAL API

As of this release, the `org.infinispan.query.dsl.Query` interface exposes the underlying Ickle query string with the `Query.getQueryString()` method. This applies to all queries regardless of whether the query was created by a DSL QueryBuilder or from a query string.

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.

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>

<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-1917 - Continuous Operations Result in RangeError

Certain clients encounter **RangeError: Index out of range** errors when performing continuous operations with JBoss Data Grid.

JDG-1914 - Errors Occur After Resolving Conflicts During Network Partition Recovery

In some cases, after a network partition occurred, JBoss Data Grid incorrectly sent broadcast messages to nodes in the merged cluster. This resulted in errors with conflict resolution.

JDG-1875 - NullPointerException in the Multimap Cache

When using the Multimap cache for some operations, the **eventType** parameter in the **convert** method of the **CacheEventConverter** interface had a null value which resulted in a **NullPointerException**.

JDG-1872 - JGroups Transport Sends Messages to All Target Nodes

The **JGroupsTransport** implementation in JBoss Data Grid sent messages all nodes in the cluster that were configured as targets. In some instances, this behavior could increase memory usage for those nodes.

JDG-1259 - Unnecessary Remote Calls When Calculating Size for Replication

When calculating the size for a replicated cache, JBoss Data Grid unnecessarily called remote nodes.

JDG-1923 - Java Method to Return Number of Processors Not Cgroup Aware

The Oracle Java method to return the number of processors available to the JVM did not take into account Control Group (Cgroup) settings for processors.

JDG-1901 - Updates to Cluster Topology Not Logged Appropriately

The JBoss Data Grid event logger did not log updates to cluster topology appropriately.

JDG-1893 - The `getQueryString` Method Not Exposed

The **Query** interface did not expose the **getQueryString** method.