



Red Hat JBoss Data Grid 7.2

7.2 Release Notes

Release Information for JBoss Data Grid

Red Hat JBoss Data Grid 7.2 7.2 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 7.2.0 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. JBOSS DATA GRID 7.2	3
1.1. INTRODUCTION TO RED HAT JBOSS DATA GRID	3
1.2. ABOUT THE RELEASE NOTES	3
CHAPTER 2. NEW FEATURES AND ENHANCEMENTS	4
2.1. DATA CONTAINER ENHANCEMENTS	4
2.2. PARTITION HANDLING ENHANCEMENTS	4
2.3. HOT ROD C# CLIENT SASL AUTHENTICATION	4
2.4. HOT ROD C++ CLIENT SASL AUTHENTICATION	4
2.5. PROTOSTREAM 4.2.0	4
2.6. HIBERNATE SEARCH 5.7.3.FINAL AND APACHE LUCENE 5.5.5	4
2.7. EAP INTEGRATION	4
2.8. SPARK CONNECTOR ENHANCEMENTS	4
2.9. SIMULTANEOUS RELEASE OF OPENSIFT IMAGES	5
2.10. CLUSTERED COUNTERS	5
2.11. CLUSTERED LOCKS	5
2.12. MULTIMAP CACHE	5
2.13. CLUSTER-WIDE DYNAMIC CACHE CREATION	5
2.14. BROADCAST QUERIES	5
2.15. CONFIGURABLE DATA FORMATS	5
CHAPTER 3. SUPPORTED CONFIGURATIONS	7
3.1. SUPPORTED CONFIGURATIONS	7
CHAPTER 4. COMPONENT VERSIONS	8
4.1. COMPONENT VERSIONS	8
CHAPTER 5. KNOWN ISSUES	9
5.1. KNOWN ISSUES	9
CHAPTER 6. RESOLVED ISSUES	11
6.1. RESOLVED ISSUES	11

CHAPTER 1. JBOSS DATA GRID 7.2

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. NEW FEATURES AND ENHANCEMENTS

2.1. DATA CONTAINER ENHANCEMENTS

JBoss Data Grid 7.2 includes many changes to the data container, including a new eviction strategy called **TinyLFU**, as well as a new storage type called **OffHeap**. The data container is configured using the new `<memory>` element, which integrates the old `<eviction>` and `<storeAsBinary>` elements, which are thus deprecated. For more information on the current data container see [Set Up JVM Memory Management](#) in the Administration and Configuration Guide.

2.2. PARTITION HANDLING ENHANCEMENTS

JBoss Data Grid 7.2 includes many updates to partition handling, including the ability to set different partition handling strategies, a new conflict manager, and several merge policies which can be used to customize how conflicts are resolved. Previously, in JBoss Data Grid, partition handling could only be turned on and off. With this update, partition handling behavior can be finely tuned. For more information on these updates see [Network Partition Recovery](#) in the Administration and Configuration Guide.

2.3. HOT ROD C# CLIENT SASL AUTHENTICATION

JBoss Data Grid 7.2 now supports SASL authentication between the C# Hot Rod client and the server. Previously, JDG only supported SASL authentication in the native Java Hot Rod Client. The C# client supports the following security mechanisms: PLAIN, DIGEST-MD5, EXTERNAL. See [Securing Interfaces](#) in the Developer Guide for more information.

2.4. HOT ROD C++ CLIENT SASL AUTHENTICATION

JBoss Data Grid 7.2 now supports SASL authentication between the C++ Hot Rod client and the server. Previously, JDG only supported SASL authentication in the native Java Hot Rod Client. The C++ client supports the following security mechanisms: PLAIN, DIGEST-MD5, EXTERNAL, GSSAPI (GSSAPI is not supported on Windows). See [Securing Interfaces](#) in the Developer Guide for more information.

2.5. PROTOSTREAM 4.2.0

JBoss Data Grid 7.2 now includes [ProtoStream](#) 4.2.0. ProtoStream is a serialization library based on Google's [Protobuf](#) data format.

2.6. HIBERNATE SEARCH 5.7.3.FINAL AND APACHE LUCENE 5.5.5

JBoss Data Grid 7.2 now includes Hibernate 5.7.3.Final and Apache Lucene 5.5.5. Both are used by JDG's Query Module to perform querying on the entire data set stored in the distributed grid. JDG can also serve as a directory provider for Hibernate Search.

2.7. EAP INTEGRATION

JBoss Data Grid 7.2 now allows caches to be configured in `standalone.xml` via a new subsystem. Using this method the cache lifecycle will not be bound to the application. The caches can also be injected via a `@Resource` annotation similar to how EAP's internal Infinispan modules are injected, just in a supported way. For more information, see [Integration with EAP](#) in the Developer Guide.

2.8. SPARK CONNECTOR ENHANCEMENTS

JBoss Data Grid 7.2 includes a new version of the Apache Spark connector which supports Spark 1.6.x and 2.x, which in turn support Scala 2.10.x, and 2.11.x. The connector provides improved configuration, simplified filtering by query String, improved Protocol Buffers support, and new support for the Dataset API, among other improvements. For more information, see [Integration with Apache Spark](#) in the Developer Guide.

2.9. SIMULTANEOUS RELEASE OF OPENSIFT IMAGES

JBoss Data Grid 7.2 will release OpenShift images for version 7.2 at the same time as the release of the standard zip archives to be used on non-containerized platforms. The OpenShift images for JBoss Data Grid 7.2 can be downloaded from the [Red Hat Container Catalog](#).

2.10. CLUSTERED COUNTERS

JBoss Data Grid 7.2 adds support for clustered counters. Clustered counters are distributed and shared across nodes in a Red Hat JBoss Data Grid cluster and allow you to record the count of objects. Clustered counters are available in Red Hat JBoss Data Grid Library Mode and Remote Client-Server Mode through the **counter** API. For more information, see [Clustered Counters](#) in the Developer Guide.

2.11. CLUSTERED LOCKS

JBoss Data Grid 7.2 adds support for clustered locks. Clustered locks are data structures that are distributed and shared across nodes in a Red Hat JBoss Data Grid cluster. Clustered locks allow you to run code that is synchronized between the nodes in a cluster. Clustered locks are available in Red Hat JBoss Data Grid Library Mode only through the **lock** API. For more information, see [Clustered Locks](#) in the Developer Guide.

2.12. MULTIMAP CACHE

JBoss Data Grid 7.2 now includes the **MultimapCache** which is a cache that maps keys to values in which each key can contain multiple values. It currently only functions in Library Mode. For more information, see [The Multimap Cache](#) in the Developer Guide.

2.13. CLUSTER-WIDE DYNAMIC CACHE CREATION

JBoss Data Grid 7.2 now supports dynamic cache creation. In Library Mode, or from a Hot Rod client, to dynamically create a cache use the **createCache()** method. For example, in Library Mode, the following method dynamically creates a cache on any new nodes that join a cluster: **manager.administration().createCache("newCacheName", "newTemplate")**. For more information, see [Library Mode](#) or [Remote Client-Server Mode](#) in the Developer Guide.

2.14. BROADCAST QUERIES

The broadcast query feature allows each node to index its own data during writes, and at query time, it sends, or "broadcasts", the query to each node. The results from each node are then combined before being returned to the caller. This is ideal for DIST caches with large indices since the amount of data transferred is the query itself and the results. For more information see [Broadcast Queries](#) in the Developer Guide.

2.15. CONFIGURABLE DATA FORMATS

JBoss Data Grid allows you to interoperate with cached data through the REST API. As of 7.2, you can configure media types that define the data format. The default media type is **application/octet-stream** with some exceptions.

In this release, JBoss Data Grid also adds support for the **application/json**, which allows you to store, retrieve, and search for JSON documents. Additionally, JBoss Data Grid provides a technology preview for querying data via the REST interface using Ickle queries in JSON format.

For more information about supported media types, see [Using the REST Interface](#) in the Developer Guide.

CHAPTER 3. SUPPORTED CONFIGURATIONS

3.1. 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 4. COMPONENT VERSIONS

4.1. 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 5. KNOWN ISSUES

5.1. 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.

JBoss Data Grid API Documentation Not Available for HotRod C++ Client in Container Installations

Issue: [JDG-1879](#)

Description: If you install the JBoss Data Grid HotRod C++ client as a container from the Red Hat Container Catalog the API documentation is not available in the `/usr/share/doc/infinispan-hotrod-cpp/api/` directory.

This issue occurs because the container image has the **nodocs** transaction flag set in the yum configuration.

Workaround: Clear the **nodocs** flag from the yum configuration and re-install the container image as follows:

```
yum --setopt=tsflags='' reinstall jdg-cpp-client
```

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 6. RESOLVED ISSUES

6.1. RESOLVED ISSUES

This release of JBoss Data Grid resolves the following issues:

JDG-1739 - Socket Connections Leak When Unknown Host Exceptions Occur

In cases where the Hot Rod client continually attempted to connect to unknown or invalid hosts, socket connections were left open. This issue was observed when the `org.infinispan` log level was set to `TRACE`.

JDG-1768 - JBoss Data Grid Expiration Values Do Not Apply to Requests from REST API

Expiration values did not apply to requests to the JBoss Data Grid REST interface if headers were not included in the requests.

JDG-1143 - Merge Function Does Not Work as Expected for the Simple Cache Interface

The `Merge` method did not work as expected in the simple Cache interface.

JDG-1127 - Error Occurs for Spatial Query Requests under High Load

An `org.hibernate.search.exception.SearchException` error would occur if multiple queries with different query parameters were run concurrently.

JDG-1039 - Hot Rod Clients Cannot Resolve Topology Addresses After Connection Failure

In dynamic environments where JBoss Data Grid nodes restart with different IP address assignments, Hot Rod clients would fail to reconnect and could not resolve the IP addresses for the nodes.

JDG-1026 - Not Possible to Set the CompatibilityProtoStreamMarshaller Class

It was not possible to set the compatibility marshaller to `CompatibilityProtoStreamMarshaller`.

JDG-1122 - Writes Via Memcached Server Result in UnknownOperationException

In some cases, when using JBoss Data Grid as a memcache service in an application, an `UnknownOperationException` exception would occur and data would not be stored when writing key value pairs of different sizes.

JDG-954 - Ickle full text range query fails with numbers

Due to a bug in `QueryRenderer`, a full text range query with numbers as the range would result in a `NullPointerException`.

JDG-989 - Cache templates with Type fail when they enable indexing

In certain cases when using indexing with cache templates, a `NullPointerException` could be thrown.

JDG-952 - Null values are counted in remote queries with aggregations

When using the Hot Rod client, if a field was annotated with `Store.YES`, null values using the field would be counted.

JDG-893 - Implement @Analyzer and @Field.analyzer for protobuf

The `@Analyzer` and `@Field.analyzer` annotations were already recognized but they were not honored.

JDG-663 - Operator '-' doesn't work in fulltext querying

When running a query like `from Entity t where t.longDescription : ('really')`, all entities that didn't contain the word `really` should have been printed, but they weren't.

JDG-664 - Right operand in term query is not being analyzed

With the default configuration and the default analyzer, queries should be "lower-cased." For example, the following two queries should return the same results: **from Entity where longDescription: 'rent'** and **from Entity where longDescription: 'RENT'**. In reality, these queries did not return the same results.