



Red Hat JBoss Data Grid 7.0

7.0.0 Release Notes

Known and resolved issues for Red Hat JBoss Data Grid 7.0.0

Red Hat JBoss Data Grid 7.0 7.0.0 Release Notes

Known and resolved issues for Red Hat JBoss Data Grid 7.0.0

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](#). 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

The Red Hat JBoss Data Grid 7.0.0 Release Notes list and provide descriptions for a series of bugzilla bugs. The bugs highlight issues that are known problems and resolved issues for the release.

Table of Contents

CHAPTER 1. INTRODUCTION TO RED HAT JBOSS DATA GRID 7.0	3
1.1. ABOUT RED HAT JBOSS DATA GRID	3
1.2. OVERVIEW	3
CHAPTER 2. NEW FEATURES AND ENHANCEMENTS	4
2.1. DISTRIBUTED STREAMS	4
2.2. REMOTE TASK EXECUTION	4
2.3. APACHE SPARK INTEGRATION	4
2.4. APACHE HADOOP INTEGRATION	4
2.5. NEW ADMINISTRATION CONSOLE FOR SERVER DEPLOYMENTS	4
2.6. CONTROLLED SHUTDOWN AND RESTART OF CLUSTER	5
2.7. NODE.JS (JAVASCRIPT) HOT ROD CLIENT	5
2.8. CASSANDRA CACHE STORE	5
2.9. HOT ROD C++ ENHANCEMENTS	5
2.10. HOT ROD C# ENHANCEMENTS	5
CHAPTER 3. DEPRECATED OR REMOVED FEATURES	6
3.1. DEPRECATED OR REMOVED FEATURES	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
APPENDIX A. REVISION HISTORY	11

CHAPTER 1. INTRODUCTION TO RED HAT JBOSS DATA GRID

7.0

Welcome to Red Hat JBoss Data Grid 7.0.0. As you become familiar with the newest version of JBoss Data Grid these Release Notes provide you with information about new features, as well as known and resolved issues. Use this document in conjunction with the entire JBoss Data Grid documentation suite, available at the Red Hat Customer Service Portal's [JBoss Data Grid documentation page](#).

[Report a bug](#)

1.1. ABOUT RED HAT JBOSS DATA GRID

Red Hat's JBoss Data Grid is an open source, distributed, in-memory key/value data store built from the Infinispan open source software project. Whether deployed in client/server mode or embedded in a Java Virtual Machine, it is built to be elastic, high performance, highly available and to scale linearly.

JBoss Data Grid is accessible for both Java and Non-Java clients. Using JBoss Data Grid, data is distributed and replicated across a manageable cluster of nodes, optionally written to disk and easily accessible using the REST, Memcached and Hot Rod protocols, or directly in process through a traditional Java Map API.

[Report a bug](#)

1.2. OVERVIEW

This document contains information about the known issues of Red Hat JBoss Data Grid version 7.0.0. Customers are requested to read this documentation prior to installing this version.

[Report a bug](#)

CHAPTER 2. NEW FEATURES AND ENHANCEMENTS

2.1. DISTRIBUTED STREAMS

Red Hat JBoss Data Grid (JDG) 7.0 introduces a distributed version of the Java 8 Stream API which enables you to perform rich analytics operations on data stored in JDG using the functional expressions available in the Stream API.

Distributed Stream is a Library mode feature, but can also be executed on JDG Server using the new Remote Task Execution feature.

[Report a bug](#)

2.2. REMOTE TASK EXECUTION

JDG 7.0 features the ability to execute tasks (business logic) on JDG Server from the Java Hot Rod client. The task can be expressed as a Java executable loaded on JDG Server or as stored JavaScript procedure which executes on the Java 8 (Nashorn) scripting engine on JDG Server.

Remote Task Execution applies to Client-Server mode.

[Report a bug](#)

2.3. APACHE SPARK INTEGRATION

JDG 7.0 introduces a Resilient Distributed Dataset (RDD) and Discretized Stream (DStream) integration with Apache Spark version 1.6. This enables you to use JDG as a highly scalable, high-performance data source for Apache Spark, executing Spark and Spark Streaming operations on data stored in JDG.

The Apache Spark integration is available in Client-Server mode.

[Report a bug](#)

2.4. APACHE HADOOP INTEGRATION

JDG 7.0 features a Hadoop InputFormat/OutputFormat integration, which enables use of JDG as a highly scalable, high performance data source for Hadoop. This enables use of tools from the Hadoop ecosystem which support InputFormat/OutputFormat for processing on data stored in JDG.

The Hadoop integration is available in Client-Server mode.

[Report a bug](#)

2.5. NEW ADMINISTRATION CONSOLE FOR SERVER DEPLOYMENTS

JDG 7.0 introduces a new Administration Console which enables you to view a JDG cluster and perform clustered operations across its nodes. Operations include creation of new caches and cache templates, starting or stopping the cluster, adding or removing nodes, and deploying or executing remote tasks.

The Administration Console is a Client-Server mode feature and requires JDG Server deployed in Domain Mode.

[Report a bug](#)

2.6. CONTROLLED SHUTDOWN AND RESTART OF CLUSTER

JDG 7.0 adds the ability to shutdown or restart a cluster in a controlled manner, with data restore from persistent storage.

This feature is available in Client-Server mode.

[Report a bug](#)

2.7. NODE.JS (JAVASCRIPT) HOT ROD CLIENT

JDG 7.0 introduces a new Node.js (JavaScript) Hot Rod client, which enables you to use JDG as a high performance distributed in-memory NoSQL store from Node.js applications.

This feature is applicable to Client-Server mode.

[Report a bug](#)

2.8. CASSANDRA CACHE STORE

JDG 7.0 introduces a new out-of-the-box Cassandra cache store, which enables you to persist the entries of a distributed cache on a shared Apache Cassandra instance.

The Cassandra cache store is available in both Library and Client-Server modes.

[Report a bug](#)

2.9. HOT ROD C++ ENHANCEMENTS

JDG 7.0 introduces asynchronous operations, querying, remote task invocation, and encryption of client/server communication using TLS/SSL, as Technical Preview features in the Hot Rod C++ client .

These enhancements are applicable to Client-Server mode.

[Report a bug](#)

2.10. HOT ROD C# ENHANCEMENTS

JDG 7.0 introduces querying, remote task invocation, and encryption of client/server communication using TLS/SSL, as Technical Preview features in the Hot Rod C# client.

These enhancements are applicable to Client-Server mode.

[Report a bug](#)

CHAPTER 3. DEPRECATED OR REMOVED FEATURES

3.1. DEPRECATED OR REMOVED FEATURES

Red Hat JBoss Data Grid (JDG) 7.0 is a new release stream. Certain features from JDG 6.6 have been removed or deprecated in JDG 7.0. For the list of such features, please refer to the JDG 7.0 Migration Guide.

[Report a bug](#)

CHAPTER 4. SUPPORTED CONFIGURATIONS

4.1. SUPPORTED CONFIGURATIONS

Supported hardware and software configurations for JDG 7.0 are available on the Customer Portal at <https://access.redhat.com/articles/2435931>.

[Report a bug](#)

CHAPTER 5. COMPONENT VERSIONS

5.1. COMPONENT VERSIONS

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

[Report a bug](#)

CHAPTER 6. KNOWN ISSUES

6.1. KNOWN ISSUES

JDG-435 - JDBC store configuration is not copied when creating new template from basic configuration

When creating a new template the JDBC cache store configuration is not copied over, and must instead be manually entered.

This is a known issue in JBoss Data Grid 7.0.0, and no workaround exists at this time.

JDG-436 - JMX operation `ClusterCacheStats.resetStatistics()` not working

When executing the JMX operation `ClusterCacheStats.resetStatistics()` none of the cluster statistics are reset.

This is a known issue in JBoss Data Grid 7.0.0, and to workaround this issue execute `Statistics.resetStatistics()` on each server individually.

JDG-474 - CLI - container and cache commands don't work in domain mode

When executing `container` or `cache` based commands through the CLI in domain mode the following message is displayed:

```
Container does not exists
```

This is a known issue in JBoss Data Grid 7.0.0, and no workaround exists at this time.

JDG-477 - The C# client does not have asynchronous versions of cache operations

Unlike the C++ and Java Hot Rod clients, the C# client does not include asynchronous versions of cache operations at this time.

This is a known issue in JBoss Data Grid 7.0.0, and a workaround exists by using the Task API to wrap the cache operation. For instance, the following example demonstrates performing an asynchronous `Put(key, value)`:

```
await Task.Run(() => cache.Put(k,v))
```

JDG-244 - Spark connector fails with CCE when compatibility mode is enabled

When attempting to insert values into a cache using the Spark connector while compatibility mode is enabled the request will fail with a `ClassCastException`.

This is a known issue in JBoss Data Grid 7.0.0, and to workaround this issue disable compatibility mode.

JDG-333 - `NotSerializableException` while executing streams via JavaScript in a cluster in DIST mode

It is not possible to use scripts that create a `Stream` in JavaScript when the cluster is in `DIST` mode. Any attempts to execute these scripts will result in a `NotSerializableException`, as the

lambdas fail when attempting to be serialized.

This is a known issue in JBoss Data Grid 7.0.0, and to workaround this issue it is recommended to manually iterator over data using an `Iterator`, or to execute lambdas after the data has been transferred from the script to the originator node.

JDG-96 - Cache events are incompatible with CDI 1.1 in JDG 7

When using CDI annotations with the CDI 1.1 API, such as the version found within JBoss EAP 7, a `java.lang.RuntimeException` is thrown; this exception is thrown as the cache events used by JBoss Data Grid 7.0 include content that has been removed by CDI 1.1.

This is a known issue in JBoss Data Grid 7.0.0. This issue is resolved in the CDI 1.2 API; to workaround this issue attach the listeners manually, as outlined in the following steps:

1. Remove CDI annotations from the listener bean, including removing all `@Observes` annotations from methods.
2. Add Listener annotations on desired methods; for example, `@CacheStarted`.
3. Manually register listeners using `AdvancedCache#addListener`.

JDG-447 - Changing transport configuration results in restarting wrong cluster

When updating one server group's transport configuration through the Administration Console and accepting the prompt to restart the server group remains running. Instead, a second server group is restarted in its place.

This is a known issue in JBoss Data Grid 7.0.0, and no workaround exists at this time.

JDG-448 - Endpoints column on cache container page shows wrong ports

When examining the port for containers in the Administration Console the `port - offset` is not considered, causing multiple containers to show at the same port.

This is a known issue in JBoss Data Grid 7.0.0, and no workaround exists at this time.

JGRP-2088 - ArrayIndexOutOfBoundsException on ClassConfigurator.get()

When attempting to use the JGroups `ENCRYPT` protocol a `ArrayIndexOutOfBoundsException` is thrown.

This is a known issue in JBoss Data Grid 7.0.0, and to workaround this issue use the `SYM_ENCRYPT` or `ASYM_ENCRYPT` protocols, as described in the JBoss Data Grid *Administration and Configuration Guide*.

[Report a bug](#)

APPENDIX A. REVISION HISTORY

Revision 7.0.0-3 Updated list of known issues. Included additional features for GA.	Sat 16 Jul 2016	Christian Huffman
Revision 7.0.0-2 Updated note on C++ Client.	Wed 27 Apr 2016	Christian Huffman
Revision 7.0.0-1 Included notes on post beta deliverables.	Wed 27 Apr 2016	Christian Huffman
Revision 7.0.0-0 Initial draft of release notes. Included new features. Included known issues.	Tue 5 Apr 2016	Christian Huffman