



Red Hat JBoss Data Grid 6

6.0.1 Release Notes

Release Notes for Red Hat JBoss Data Grid 6.0.1

Edition 1

Red Hat JBoss Data Grid 6 6.0.1 Release Notes

Release Notes for Red Hat JBoss Data Grid 6.0.1
Edition 1

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

Legal Notice

Copyright © 2012 .

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 6.0.1 Release Notes list and provide descriptions for a series of bugzilla bugs. The bugs highlight either issues that are known problems for the relevant release or bugs that have now been resolved.

Table of Contents

CHAPTER 1. ABOUT JBOSS DATA GRID 6.0.1	3
CHAPTER 2. KNOWN ISSUES	4
CHAPTER 3. RESOLVED ISSUES	7
APPENDIX A. REVISION HISTORY	9

CHAPTER 1. ABOUT JBOSS DATA GRID 6.0.1

Red Hat's JBoss Data Grid is an open source, distributed, in-memory key/value datastore built from the Infinispan open source software project. Whether deployed in client/server mode or embedded in a Java VM, it is built to be elastic, performant, 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.

Both Remote Client-Server Mode and Library Mode are fully supported in JBoss Data Grid 6.0.1.

CHAPTER 2. KNOWN ISSUES

The following issues are known to exist in JBoss Data Grid 6.0.1 and will be fixed in a subsequent release.

BZ#745865 - ServerWorker thread naming problem

If a server is used for both Hot Rod and Memcached, it is not possible to distinguish between the worker threads for each protocol because they are all named "MemcachedServerWorker". This does not affect the functionality of the server.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#760895 - Reopened: Error detecting crashed member during shutdown of EDG 6.0.0.Beta

Occasionally, when shutting down nodes in a cluster, the following message is reported: "ERROR [org.infinispan.server.hotrod.HotRodServer] ISPN006002: Error detecting crashed member: java.lang.IllegalStateException: Cache '___hotRodTopologyCache' is in 'STOPPING' state and this is an invocation not belonging to an on-going transaction, so it does not accept new invocations. Either restart it or recreate the cache container."

This is due to the fact that a node has detected another node's shutdown and is attempting to update the topology cache while itself is also shutting down. The message is harmless, and it will be removed in a future release

This behavior persists in JBoss Data Grid 6.0.1.

BZ#807674 - JDBC Cache Stores using a JTA Data Source do not participate in cache transactions

In JBoss Data Grid's library mode, JDBC cache stores can be configured to use a JTA-aware datasource. However operations performed on a cache backed by such a store during a JTA transaction, will be persisted to the store outside of the transaction's scope. This issue is not applicable to JBoss Data Grid's Remote Client-Server mode because all cache operations are non-transactional.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#807741 - Reopened: Invalid magic number

The issue is a race condition on the Hot Rod server, which can lead to topologies being sent erroneously as a result of addition of a new node to the cluster. When the issue appears, clients view "Invalid magic number" error messages due to unexpected data within the stream.

In such a situation, the recommended approach is to restart the client. If the client is not restarted, the client may recover after the unexpected data is consumed but this is not guaranteed. If the client recovers with a restart, the view topology displayed does not display one of the nodes added, resulting in uneven request distribution.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#808623 - Some entries not available during view change

In rare circumstances, when a node leaves the cluster, instead of going directly to a new cluster view that displays all nodes save the note that has departed, the cluster splits into two partitions which then merge after a short amount of time. During this time, some nodes do not have access to all the data that previously existed in the cache. After the merge, all nodes regain access to all the data, but changes made during the split may be lost or be visible only to a part of the cluster.

Normally, when the view changes because a node joins or leaves, the cache data is rebalanced on the new cluster members. However, if the number of nodes that leaves the cluster in quick succession equals or is greater than the value of `numOwners`, keys for the departed nodes are lost. This occurs during a network split as well - regardless of the reasons for the partitions forming, at least one partition will not have all the data (assuming cluster size is greater than `numOwners`).

While there are multiple partitions, each one can make changes to the data independently, so a remote client will see inconsistencies in the data. When merging, JBoss Data Grid does not attempt to resolve these inconsistencies, so different nodes may hold different values even after the merge.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#818092 - NPE in Externalizer on shutdown

A workaround has been added to avoid the `NullPointerException` while AS7 implements the right order to stop Infinispan services.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#818863 - Reopened: Tests for UNORDERED and LRU eviction strategies fail on IBM JDK

The `LinkedHashMap` implementation in IBM's JDK sometimes behaves erratically when extended (as is done by the eviction strategy code). This wrong behaviour is exposed by the JBoss Data Grid Test Suite. The recommendation is, if using eviction, to use other JDKs (Oracle or OpenJDK) which are not affected by this issue.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#822815 - NPE during JGroups Channel Service startup

Occasionally, when starting a JBoss Data Grid server, the JGroups subsystem would not start because of a `NullPointerException` during service installation, leaving the server in an unusable state. This situation does not affect data integrity within the cluster, and simply killing the server and restarting it solves the problem.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#829813 - Command line bind address being ignored for HTTP and Infinispan servers

The JBoss Data Grid server by default binds its listening ports to a loopback address (`127.0.0.1`). The `-b` switch can be used to modify the address on which to bind the public interface. However, since the JBoss Data Grid endpoints are bound to the management interface for security reasons, the `-b` switch does not affect them. The user should modify the `standalone.xml` configuration file to place the endpoints on the public interface:

```
<socket-binding name="hotrod" interface="public" port="11222"/>
```

After the above modification, the `-b` switch will determine the network address on which the hot rod port will be bound.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#841889 - Transaction leak caused by reordering between prepare and commit

This may cause locks to not be released and as a result, other transactions may not be able to write that data (though they are able to read it). As a workaround, one can enable transaction recovery, then use the JMX recovery hooks in order to cleanup the pending lock.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#841891 - Potential tx lock leaks when nodes are added to the cluster

This situation only occurs during topology changes and results in certain keys remaining locked after a transaction is finished. This means that no other transaction is able to write to the given key while it is locked (though it can still be read). As a workaround, one can enable transaction recovery and then use the JMX recovery hooks in order to cleanup the pending lock.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#847062 - Pre-Invocation flag PUT_FOR_EXTERNAL_READ throws exception

Users should not pass in Flag.PUT_FOR_EXTERNAL_READ as it is designed for internal use only. Alternatively, call Cache.putForExternalRead() operation.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#847809 - Cluster with non-shared JDBC cache store has too many entries after node failure

The root cause of this problem still needs to be analyzed. As a workaround, use a shared cache store instead of local cache stores.

This behavior persists in JBoss Data Grid 6.0.1.

BZ#854665 - Coordinator tries to install new view after graceful shutdown

In JBoss Data Grid 6.0.0, when gracefully stopping a coordinator node, the node itself would log an attempt to install a new clustering view which would fail. This is harmless, since the new coordinator would in fact perform the proper view installation.

This behavior persists in JBoss Data Grid 6.0.1.

CHAPTER 3. RESOLVED ISSUES

The following issues have been resolved in JBoss Data Grid 6.0.1.

BZ#810155 - Default configuration may not be optimal for a 2-node cluster

A sample configuration is now included in the JBoss Data Grid Server distribution called `docs/examples/configs/standalone-two-nodes.xml`, which provides better performance when using clusters formed by only two nodes. This configuration changes the default JGroups configuration to use the UNICAST protocol instead of the UNICAST2 and RSVP combination of protocols. For larger clusters the UNICAST2 and RSVP combination is recommended.

BZ#841887 - Transactional listeners method order problem

In JBoss Data Grid 6.0.0 listeners on remote nodes would be invoked in a different order compared the originator node. This has now been resolved and remote nodes will invoke listeners in the same, correct order as on the originator nodes.

BZ#841888 - Make connection refused exceptions TRACE

In JBoss Data Grid 6.0.0 the HotRod connection pooling check for idle connections would log spurious WARN messages about not being able to create a connection. These messages are now logged at the TRACE level since they are harmless and can safely be ignored.

BZ#841890 - Duplicate counts for some stats

In JBoss Data Grid 6.0.0 some statistics were counted twice. This has now been corrected and all statistics are counted once as expected.

BZ#841892 - ping_on_startup ignored

In JBoss Data Grid 6.0.0 the HotRod client option `infinispan.client.hotrod.ping_on_startup` was ignored and the client pinged the servers on start up. This is now resolved and the client's behavior is in accordance with the `infinispan.client.hotrod.ping_on_startup` option.

BZ#841893 - Ping should try all servers in server list.

In JBoss Data Grid 6.0.0 HotRod clients attempting to retrieve topology information using the protocol's PING message from a non-functional server would fail immediately without retrying. This issue has been resolved and the clients will retry the failed PING operation.

BZ#841896 - JDBC cache store should quote generated table name

In JBoss Data Grid 6.0.0 the JDBC cache store would not enclose generated table names within quotes, which caused problems when attempting to use special characters. This has now been resolved, and table names are quoted.

BZ#846729 - Some DefaultCacheManager constructors still performing legacy configuration adaptation

In JBoss Data Grid 6.0.0 some of the `org.infinispan.manager.DefaultCacheManager` constructors using the `org.infinispan.configuration` API would invoke the adapter to convert configuration to the legacy `org.infinispan.config` API. This would cause some properties, such as a user-specified classloader, to be lost. This has now been resolved and the legacy adapter is no longer invoked in these situations.

BZ#846731 - AtomicHashMapProxy calls Cache.getCacheConfiguration() triggering legacy config adaptation.

In JBoss Data Grid 6.0.0 the AtomicHashMapProxy class would trigger an invocation to the legacy configuration adapter causing performance degradation and a spurious INFO log message. This is now resolved.

BZ#847711 - Cache view installation takes too long

When the coordinator shuts down, it tries to install a new cache view that does not include itself. Because there are multiple nodes shutting down at the same time, two or more nodes can see themselves as the coordinator and will try to install conflicting views, failing and repeating the process for a very long time. This is now resolved by automatically rejecting outside cache views installation requests on the current JGroups coordinator. This way each one of the "coordinators" can shut down without interfering with the others.

BZ#854662 - Upgrade to JGroups 3.0.13.Final

The version of JGroups bundled with JBoss Data Grid 6.0.1 is upgraded to 3.0.13.Final which provides bug fixes and performance improvements.

APPENDIX A. REVISION HISTORY

Revision 0.0-9.400 Rebuild with publican 4.0.0	2013-10-31	Rüdiger Landmann
Revision 0.0-9 Published with updated product name.	Mon Aug 05 2013	Misha Husnain Ali
Revision 0.0-8 Updating with doc text editing.	Tue Sept 18 2012	Gemma Sheldon
Revision 0.0-7 Rebrew to remove BZ#856101 as this is not a bug.	Fri Sept 14 2012	Gemma Sheldon
Revision 0.0-6.1 Final brew updating bug statuses.	Thur Sept 13 2012	Gemma Sheldon
Revision 0.0-5 Updated with 25 bugs in total for release notes. Bugs still ON QA are highlighted. Updated About JBoss Data Grid.	Wed Sept 12 2012	Gemma Sheldon
Revision 0.0-4 Updated with highlighting for bugs currently ON_QA pending approval for their inclusion in the Release Notes. BZ#847711 and BZ#854662 in resolved section for now.	Mon Sept 10 2012	Gemma Sheldon
Revision 0.0-3 Corrected typos in 1 bug.	Fri Sept 7 2012	Gemma Sheldon
Revision 0.0-2 Updated to include CCFR information for 6.0.1 Release Notes.	Fri Sept 7 2012	Gemma Sheldon
Revision 0.0-1 Initial creation by publican	Thur Sept 6 2012	Gemma Sheldon