



Red Hat JBoss Data Grid 7.0

Migration Guide

For use with Red Hat JBoss Data Grid 7.0

Red Hat JBoss Data Grid 7.0 Migration Guide

For use with Red Hat JBoss Data Grid 7.0

Misha Husnain Ali
Red Hat Engineering Content Services
mhusnain@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](https://creativecommons.org/licenses/by-sa/3.0/). 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

This guide presents information about migrating to JBoss Data Grid 7.0

Table of Contents

CHAPTER 1. INTRODUCTION	3
1.1. ABOUT THE MIGRATION GUIDE	3
CHAPTER 2. PRODUCT CHANGES	4
2.1. LIBRARY MODE CONFIGURATION UPDATE	4
2.2. CLI COMMANDS	5
2.3. SERVER MANAGMENT	5
2.4. JGROUPS ENCRYPT PROTOCOL DEPRECATED	5
CHAPTER 3. APPLICATION CHANGES	6
3.1. NOTIFYINGFUTURE REPLACED BY COMPLETABLEFUTURE	6
3.2. JBOSS MODULES CDI DEPENDENCY	6
3.3. SPRING 3 SUPPORT REMOVED	6
3.4. REMOTECACHE DEPRECATED METHODS	6
3.5. INTERCEPTORS DEPRECATED	7
3.6. MAP/REDUCE REPLACED BY DISTRIBUTED STREAMS	7
3.7. HIBERNATE SEARCH	7
APPENDIX A. REVISION HISTORY	11

CHAPTER 1. INTRODUCTION

Red Hat JBoss Data Grid is a distributed in-memory data grid, which provides the following capabilities:

- Schemaless key-value store – JBoss Data Grid is a NoSQL database that provides the flexibility to store different objects without a fixed data model.
- Grid-based data storage – JBoss Data Grid is designed to easily replicate data across multiple nodes.
- Elastic scaling – Adding and removing nodes is simple and non-disruptive.
- Multiple access protocols – It is easy to access the data grid using REST, Memcached, Hot Rod, or simple map-like API.

[Report a bug](#)

1.1. ABOUT THE MIGRATION GUIDE

The purpose of this guide is to document the changes that are required to successfully run and deploy Red Hat JBoss Data Grid 6 applications on Red Hat JBoss Data Grid 7. It provides information on how to resolve deployment and runtime problems and to how prevent changes in application behavior. This is the first step in moving to the new platform. Once the application is successfully deployed and running, plans can be made to upgrade individual components to use the new functions and features of JBoss Data Grid 7.

[Report a bug](#)

CHAPTER 2. PRODUCT CHANGES

2.1. LIBRARY MODE CONFIGURATION UPDATE

The XML schema for Library mode configurations has been changed to be more consistent with the remote client-server schema. This results in several changes that must be made to the configuration file.

To assist in these configurations a tool is provided in the library distribution. The following steps discuss using this tool to automatically convert existing infinispn configuration files to the latest supported schema.

Procedure 2.1. Windows Instructions for Using the Library Mode Configuration Converter

1. Open a command prompt.
2. Navigate to the **bin** directory of the library distribution:

```
cd C:\path\to\jboss-datagrid-${jdg-version}-library\bin\
```

3. If any custom classes were referenced in the original configuration then the .jar files containing these classes must be placed in the **C:\path\to\jboss-datagrid-\${jdg-version}-library** directory. If no custom classes were referenced then this step may be skipped.
4. Execute the **config-converter.bat** utility, passing in the name of the existing **infinispn_6.0.xml** configuration file and the destination as arguments:

```
config-converter C:\path\to\infinispn_6.0.xml  
C:\path\to\infinispn-new.xml
```

5. Ensure that **infinispn-new.xml** has been created with the new Library mode configurations. This file may now be used in Library mode applications with JBoss Data Grid 7.0.

Procedure 2.2. Linux Instructions for Using the Library Mode Configuration Converter

1. Open a console to execute the following commands.
2. Navigate to the **bin/** directory of the library distribution:

```
cd /path/to/jboss-datagrid-${jdg-version}-library/bin/
```

3. If any custom classes were referenced in the original configuration then the .jar files containing these classes must be placed in the **/path/to/jboss-datagrid-\${jdg-version}-library/** directory. If no custom classes were referenced then this step may be skipped.
4. Execute the **config-converter.sh** utility, passing in the name of the existing **infinispn_6.0.xml** configuration file and the destination as arguments:

```
./config-converter.sh /path/to/infinispn_6.0.xml  
/path/to/infinispn-new.xml
```

5. Ensure that **infinispn-new.xml** has been created with the new Library mode configurations. This file may now be used in Library mode applications with JBoss Data Grid 7.0.



NOTE

If no destination file is specified, such as **infinispan-new.xml** in the above examples, the utility will instead display the converted xml to the console.

[Report a bug](#)

2.2. CLI COMMANDS

The following CLI commands have been changed for JBoss Data Grid 7.0:

- **clear** - This command has been replaced by **clearcache**.
- **help \$COMMAND** - Instead of executing **help** per command, each command now has a built in **--help** flag that is used instead.

[Report a bug](#)

2.3. SERVER MANAGMENT

The default JBoss Data Grid server management port has been changed to **9990**.

[Report a bug](#)

2.4. JGROUPS ENCRYPT PROTOCOL DEPRECATED

The JGroups **ENCRYPT** protocol has been deprecated, and should be replaced with either the **SYM_ENCRYPT** or **ASYM_ENCRYPT** protocols. Both of these protocols must be placed directly under **NAKACK2**. Configuration details are provided in the *Administration and Configuration Guide*; however, example configurations for each protocol are shown below:

Example 2.1. SYM_ENCRYPT Configuration Example

```
<SYM_ENCRYPT sym_algorithm="AES"
    encrypt_entire_message="true"
    keystore_name="defaultStore.keystore"
    store_password="changeit"
    alias="myKey"/>
```

Example 2.2. ASYM_ENCRYPT Configuration Example

```
<ASYM_ENCRYPT encrypt_entire_message="true"
    sym_keylength="128"
    sym_algorithm="AES/ECB/PKCS5Padding"
    asym_keylength="512"
    asym_algorithm="RSA"/>
```

[Report a bug](#)

CHAPTER 3. APPLICATION CHANGES

3.1. NOTIFYINGFUTURE REPLACED BY COMPLETABLEFUTURE

The **NotifyingFuture** interface has been removed in JBoss Data Grid 7.0 and replaced with **CompletableFuture**. This class was replaced due to **CompletableFuture** being included as a standard class in Java 8, allowing for additional features such as chaining, implementing constructs such as conditionals, and timeouts to be utilized without introducing additional custom code.

Due to this change any applications utilizing **NotifyingFuture** will need to be replaced with **CompletableFuture** to function correctly with JBoss Data Grid 7.0.

[Report a bug](#)

3.2. JBOSS MODULES CDI DEPENDENCY

JBoss Modules CDI Dependencies

The package containing the CDI classes has changed. All of these classes have moved from **org.infinispan.cdi** to **org.infinispan.cdi.embedded**, and any applications utilizing CDI should have their imports updated to use the new packages.

Configuration of CDI/JCache

CDI Extensions have been split into two separate types of classes:

1. CDI Classes that begin with **Injected** before the classname, such as **org.infinispan.jcache.annotation.InjectedExceptionInterceptor**. These classes are intended for use in managed environments, such as JBoss EAP.
2. CDI classes without **Injected** before the classname, such as **org.infinispan.jcache.annotation.CacheExceptionHandler** - This class should be used in standalone environments.

[Report a bug](#)

3.3. SPRING 3 SUPPORT REMOVED

Spring 3 support has been removed in JBoss Data Grid 7.0. Any applications will need to be rewritten to use Spring 4.

[Report a bug](#)

3.4. REMOTECACHE DEPRECATED METHODS

The following methods have been deprecated in JBoss Data Grid 7.0:

- **remoteCache.getVersioned** - This method has been replaced by **remoteCache.getWithMetadata(Object)**.
- **remoteCache.getBulk**, **getBulk(size)** - Both of these methods have proven expensive for large data sets. The **remoteCache.retrieveEntries*** methods offer a lazy, pull-style, approach which retrieves bulk data more efficiently.

[Report a bug](#)

3.5. INTERCEPTORS DEPRECATED

Support for custom interceptors is being deprecated in JBoss Data Grid 7.0. A new method of executing custom interceptors is expected to be introduced in JBoss Data Grid 7.1. In addition, the interceptor stack is part of JBoss Data Grid's internal API, and is subject to change from release to release. Due to this it is not recommended to use custom interceptors directly from your application.

[Report a bug](#)

3.6. MAP/REDUCE REPLACED BY DISTRIBUTED STREAMS

Map/Reduce has been removed in JBoss Data Grid 7.0.0. This feature has been replaced by Distributed Streams, which were shown to have better performance.

[Report a bug](#)

3.7. HIBERNATE SEARCH

The underlying Hibernate Search modules have been upgraded from 4.6 to 5.6, and as such there have been significant changes to this component. The following sections walk through all updates necessary to use the new version.



IMPORTANT

As the Hibernate Search modules have been updated, it is necessary to delete all existing indexes and rebuild the index, as indexes produced by previous versions are not compatible with the new version in use by JBoss Data Grid 7.0. Instructions for rebuilding the index are found in the *Developer Guide*.

[Report a bug](#)

3.7.1. Hibernate Search Class Name and Package Updates

Several classes have been relocated to a different package than their previous package; the following table includes the original class name and package along with the new location.

Table 3.1. Hibernate Search Class Package Updates

JDG 6.x Location	JDG 7.0 Location
org.hibernate.search.engine.impl.SearchMappingBuilder	org.hibernate.search.engine.spi.SearchMappingBuilder
org.hibernate.search.Environment	org.hibernate.search.cfg.Environment
org.hibernate.search.FullTextFilter	org.hibernate.search.filter.FullTextFilter
org.hibernate.search.indexes.impl.DirectoryBasedIndexManager	org.hibernate.search.indexes.spi.DirectoryBasedIndexManager

JDG 6.x Location	JDG 7.0 Location
org.hibernate.search.infinispan.impl.InfinispanDirectoryProvider	org.hibernate.search.infinispan.spi.InfinispanDirectoryProvider
org.hibernate.search.ProjectionConstants	org.hibernate.search.engine.ProjectionConstants
org.hibernate.search.SearchException	org.hibernate.search.exception.SearchException
org.hibernate.search.spi.MassIndexerFactory	org.hibernate.search.batchindexing.spi.MassIndexerFactory
org.hibernate.search.spi.SearchFactoryBuilder	org.hibernate.search.spi.SearchIntegratorBuilder
org.hibernate.search.spi.SearchFactoryIntegrator	org.hibernate.search.spi.SearchIntegrator
org.hibernate.search.Version	org.hibernate.search.engine.Version
org.apache.lucene.queryParser.QueryParser	org.apache.lucene.queryparser.classic.QueryParser

Table 3.2. Hibernate Search Class Name Updates

JDG 6.x Name	JDG 7.0 Name
SpatialMode.GRID	SpatialMode.HASH

[Report a bug](#)

3.7.2. Hibernate Search Additions

The following additions have been made to the Hibernate Search modules:

- A new method, **getShardIdentifiersForDeletion()**, has been added to **org.hibernate.search.store.ShardIdentifierProvider**. Any existing implementations which were not derived from **ShardIdentifierProviderTemplate** should be updated.
- The Faceting engine has been significantly updated; however, a **@Facet** or **@Facets** annotation must be used on any fields intended for faceting.

[Report a bug](#)

3.7.3. Hibernate Search Lucene Updates

Number and Date Index Format

Numbers and dates are now indexed as numeric fields by default. Properties of type **Date**, **Calendar**, **int**, **long**, **float**, **double**, and their corresponding wrappers, are no longer indexed as strings; instead, they use Lucene's appropriate numeric encoding. If any query previously targeted a string

encoded field, but is not encoded numerically, it will need to be updated. Numeric fields must be searched with a **NumericRangeQuery**; however, if the Search query DSL is in use it should generate the correct query. In addition, any fields targeted by faceting should now be string encoded.



NOTE

Date and **Calendar** instances are encoded as a **long** value, representing the number of milliseconds since **January 1, 1970, 00:00:00 GMT**.

Any **id** fields are an exception to this rule. Even when these fields are represented by a numeric type they will continue to be indexed as a string keyword by default.

The changes mean the use of **@NumericField** is now obsolete, unless a custom precision for the numeric encoding is explicitly specified.

Previous, string based, index format may be used by specifying a string encoding field bridge. For integers this is **org.hibernate.search.bridge.builtin.IntegerBridge**; other publicly available field bridges may be found in the **org.hibernate.search.bridge.builtin** package.

For dates and calendars the indexing format may be changed via the **EncodingType** enum; for example, **@DateBridge(encoding=EncodingType.STRING)** or **@CalendarBridge(encoding=EncodingType.STRING)**.

Id fields of embedded relations are no longer indexed by default

When using an **@IndexedEmbedded** annotation to include fields from a related entity, previously the **id** of the related entity was included. These fields are no longer included by default, but may be enabled by using the **includeEmbeddedObjectId** attribute of the **@IndexedEmbedded** annotation, as seen below:

```
@IndexedEmbedded(includeEmbeddedObjectId=true)
```

Sorting Query Results

The **SortField** API has been modified to use **SortField.Type.\${CLASSNAME}** instead of the previous **SortField.\${CLASSNAME}**. This behavior is highlighted below:

```
// Previous API call
fulltextQuery.setSort( new Sort( new SortField( "title", SortField.STRING
) ) );

// New API call
fulltextQuery.setSort( new Sort( new SortField( "title",
SortField.Type.STRING ) ) );
```

Any applications not precisely matching the Lucene field encoding for sorting options in Lucene 4 would have worked without error; however, the new version of Lucene 5 now correctly catches these errors and will report a runtime exception. The expected field encoding must be matched according to its type. A field indexed as a **String** requires a single term, even after analysis, a numeric field encoded as an Integer must match **SortField.Type.INT**, a numeric field encoded as a Long must be sorted with **SortField.Type.LONG**, and so forth.

Faceting on Numeric Fields

Similar to Sorting Query Results, each target field must match an appropriate kind of Faceting Query. For instance, a Numeric Facet Query must match a field that was indexed exclusively via the matching Numeric type.

Null value tokens for Numeric Fields

When using `@Field(indexNullAs=)` to encode a null marker value in the index, the type of marker must be compatible with all other values which are indexed in the same field. Previously a string keyword, such as `null`, could be encoded with numeric fields; however, this is no longer allowed. A number must now be used to represent the null value, such as `-1`.

Lucene Classes Moved

Many of the **Analyzer**, **TokenFilter**, and **CharFilter** classes have been moved, and it is recommended to examine [Lucene's API documentation](#) to find the correct replacement.

Similarly, many of the Solr utility classes, such as **TokenizerFactory** or **TokenFilterFactory** were moved into Apache Lucene. Any applications using these classes should also be updated using the [Lucene's API documentation](#).

Due to this refactoring the Solr dependencies are now incorporated within Lucene, and no longer need to be specified; any applications that previously had Solr dependencies defined should remove them and only use the Lucene dependencies.

[Report a bug](#)

3.7.4. Hibernate Search Deprecated Classes

The following list are deprecated classes within the Hibernate Search modules:

- **org.hibernate.search.annotations.Key** - there is no replacement for this class; however, it is no longer required for users to provide a **Key** object.
- **ContainedInMapping#numericField()** - Use **field().numericField()** instead.

[Report a bug](#)

APPENDIX A. REVISION HISTORY

Revision 7.0.0-1	Mon 18 Jul 2016	Christian Huffman
Included application changes. Added JGroups ENCRYPT protocol information.		
Revision 7.0.0-0	Wed 3 Feb 2016	Christian Huffman
Created initial draft.		