



# Red Hat JBoss Data Virtualization 6.3 Migration Guide

---

This guide is for installation teams

Red Hat Customer Content  
Services



# Red Hat JBoss Data Virtualization 6.3 Migration Guide

---

This guide is for installation teams

Red Hat Customer Content Services

## 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 document guides users through the process for migrating to Red Hat JBoss Data Virtualization 6 from previous versions.

---

## Table of Contents

<b>Chapter 1. Read Me</b> .....	<b>2</b>
1.1. Back Up Your Data	2
1.2. Variable Name: EAP_HOME	2
1.3. Variable Name: MODE	2
1.4. Red Hat Documentation Site	2
<b>Chapter 2. Migration</b> .....	<b>3</b>
2.1. Product Name Change	3
2.2. Migrating Red Hat JBoss Data Virtualization	3
2.3. Migrating Hierarchical Data	4
2.4. LDAP Compliance	4
<b>Chapter 3. Changes</b> .....	<b>6</b>
3.1. Changes to Red Hat JBoss Data Virtualization	6
3.2. Changes to the Hierarchical Database System	6
<b>Chapter 4. Backup and Restore</b> .....	<b>10</b>
4.1. Migrating from a Previous Release	10
4.2. The Repository Manager	10
4.3. Backup a Repository	11
4.4. Restore a Repository	11
<b>Chapter 5. Supported Configurations</b> .....	<b>13</b>
5.1. Supported Data Sources and Translators	13
<b>Appendix A. Revision History</b> .....	<b>15</b>

## Chapter 1. Read Me

### 1.1. Back Up Your Data



#### Warning

Red Hat recommends that you back up your system settings and data before undertaking any of the configuration tasks mentioned in this book.

### 1.2. Variable Name: EAP\_HOME

**EAP\_HOME** refers to the root directory of the Red Hat JBoss Enterprise Application Platform installation on which JBoss Data Virtualization has been deployed.

### 1.3. Variable Name: MODE

**MODE** will either be **standalone** or **domain** depending on whether JBoss Data Virtualization is running in standalone or domain mode. Substitute one of these whenever you see **MODE** in a file path in this documentation. (You need to set this variable yourself, based on where the product has been installed in your directory structure.)

### 1.4. Red Hat Documentation Site

Red Hat's official documentation site is available at <https://access.redhat.com/site/documentation/>. There you will find the latest version of every book, including this one.

## Chapter 2. Migration

### 2.1. Product Name Change



#### Note

This product was previously called JBoss Enterprise Data Services, or EDS.

### 2.2. Migrating Red Hat JBoss Data Virtualization

This topic lists tasks that you might have to implement in order to migrate to Red Hat JBoss Data Virtualization 6.

#### What will I have to do?

- To migrate to Red Hat JBoss Data Virtualization 6, users will need to install a fresh instance to redeploy customer applications and artifacts.
- If users are migrating from a version prior to Red Hat JBoss Enterprise Data Services 5.3.1, then the REST WARs will need to be regenerated due to changes in CXF. Otherwise the WAR will not deploy due to the WSDL not being found.
- Previously, null values were ranked low by default. This could have a negative impact on performance and is not needed in many situations. The ranking of null values are no longer pushed down by default. Users will have to set the `org.teiid.pushdownDefaultNullOrder` property to preserve the old behavior if required.



#### Note

- The old update procedure syntax is no longer supported. Therefore, users may have to revalidate models and rebuild the VDB.
- Users may have to reimport VDBs stored by the hierarchical database in order to update node types.
- Any custom translator or connector (resource adapter) will require migrating. It will need to be recompiled and new configuration and packaging will be needed. This is due to the changes from EAP 5 to EAP 6. Also, the translator documentation should be read, because additional features have been added (for example, DDL support), that would require additional updates to the translator.

#### How will migration impact modeling?

- Users may have issues when validating transformations after importing a project set that were built in the previous version of Teiid Designer. Users may have to adjust transformations in their models to account for this.



### Note

When you migrate models containing REST procedures from 7.7.x version, edit the XMI file and change the key names. The designer tool then displays the value for the given properties.

```
key= rest:uri is now key= REST:URI  
And  
key= rest:restMethod is now key= REST:METHOD
```

- » When importing any of the EDS project, import all the `.xmi` files one-by-one. This helps you to eliminate all the possible validation errors when importing the complete project set. However, do not import the `.project` file.



### Warning

Do not remove old data before verifying the migration to JBoss Data Virtualization 6 was successful.

## 2.3. Migrating Hierarchical Data

Data stored in the provided hierarchical database can be migrated from previous versions of JBoss Data Virtualization using the new backup and restore functions.

To migrate a repository, back it up from your previous installation, then restore it in your new JBoss Data Virtualization 6 installation.

Each backup contains all content in all workspaces for a single repository. Therefore, to migrate to release 6, the backup/restore process will need to be followed for each repository.



### Note

A migration tool will be provided to backup from previous installations, compatible with the new backup and restore functions.



### Warning

Do not remove old data before verifying the migration to JBoss Data Virtualization 6 is successful.

#### See Also:

- » [Section 4.1, “Migrating from a Previous Release”](#)

## 2.4. LDAP Compliance



It is recommended that customers who have utilized the internal JDBC membership domain from releases prior to MetaMatrix 5.5 migrate those users and groups to an LDAP compliant directory server.

Refer to the *Red Hat JBoss Enterprise Application Platform Security Guide* for directions on using an LDAP directory server. Please contact technical support if you require additional guidance in the migration process.

Several free and open source directory servers include:

- ✦ The Fedora Directory Server - <http://directory.fedoraproject.org/>
- ✦ Open LDAP - <http://www.openldap.org/>
- ✦ Apache Directory Server - <http://directory.apache.org/>

## Chapter 3. Changes

### 3.1. Changes to Red Hat JBoss Data Virtualization

#### Installation

- ✦ In 5.x releases, Red Hat JBoss Data Virtualization depended on the Red Hat JBoss Services Oriented Architecture Platform (SOA-P). In this release, Red Hat JBoss Data Virtualization is installed on the Red Hat JBoss Enterprise Application Platform (EAP) directly. Users will be able to install on top of an existing JBoss EAP installation or with a fresh installation.
- ✦ In 5.xs releases, Red Hat JBoss Data Virtualization could operate under a number of *profiles*. In this release, Red Hat JBoss Data Virtualization can be installed in one of the two JBoss EAP *modes*: standalone mode (default) or domain mode (for clustering).

#### Quickstarts

Quickstarts now rely on Maven instead of Apache Ant.

#### Tooling Support

Teiid Designer and ModeShape Tools plugins for Red Hat JBoss Developer Studio will support both version 5 and 6 of JBoss Data Virtualization.

#### Week Function

Red Hat JBoss Data Virtualization now consistently returns the ISO 8601 standard for WEEK. Previously, "locale" had been used by default for many databases. This could potentially break functionality for users of databases such as h2, SQL Server and Sybase.

#### ibm.jdk Dependency Missing from org.jboss.security.negotiation Module

The `ibm.jdk` dependency is not added to the `org.jboss.security.negotiation` module, including the patched version in the `.overlay` directory. As this is not fixed in Red Hat JBoss EAP, you will have to manually add the dependency every time you patch EAP if you want to make use of it.

### 3.2. Changes to the Hierarchical Database System

#### APIs

For both releases 5 and 6, client applications use the standard JCR 2.0 API to interact with repositories. So most (if not all) of your application code can remain unchanged.

A small API (`org.modeshape.jcr.api`) was provided in release 5 that extended the standard JCR API with a few additional capabilities. This release supports and slightly expands on this API. Several of the methods and interfaces in the API were deprecated in the previous release. These have now been removed.

The URLs for the RESTful API have changed. This release provides an improved RESTful API which is now the default. The RESTful client library is compatible with the database system provided in both versions 5 and 6.

The WebDAV API has been improved and the JDBC driver is also still supported.

#### Caching

In release 5, a single Service Provider Interface (SPI) was used for both storage and access connectors. However, this has been changed as part of a move to a more scalable caching

connectors. However, this has been changed as part of a move to a more scalable caching solution.

In this release, each repository uses a single cache for both its caching and storage system. All workspace content (except for binary value storage) and system content is stored within this single cache.

## Federation

In this release, the JCR API can be used to access information in external systems. The concept of a *connector* has been reintroduced as a mechanism to achieve this. Repositories will no longer consist entirely of federated content. Instead, every repository will store its own content, but you will be able to federate and integrate content from external systems into the repository.



### Note

Technically, even in the previous release a storage connector was required to store the repository's system content, so it was never actually possible to have a repository that consisted entirely of federated content.

## Binary Storage

Repositories can now handle extremely large binary values with several storage options:

- ✦ Binary values can be stored on the local file system. This can be a regular directory, a network share or even a temporary directory. This option is generally fast, safe (as safe as your file system), and native locks are used to prevent multiple processes from conflicting. It is recommended for local (non-clustered), embedded repositories.
- ✦ Binary values can be stored in cache. Although it is possible to use the same cache as the rest of the repository, it is often better to use two other caches and to configure those caches specifically for what they store. For example, one cache is used to store metadata about the binary values; this metadata is small, lightweight, and can thus be replicated across the cluster. The other cache is used to store the actual binary data, separated into chunks (usually up to 1MB in size), and for this type of cache, distributing the data across the cluster is often desirable.
- ✦ Binary values can be stored in a relational database. This is recommended only when you intend to persist all content inside a relational database. The binary data is broken into chunks (usually up to 1MB in size).
- ✦ Binary values can be stored in MongoDB. This storage option has not been thoroughly tested, but can be considered as an option.
- ✦ Binary values can be stored in a custom store. A Service Provider Interface (SPI) is now provided for implementing your own binary storage.

## Sequencers

Each repository is now configured with its own sequencers. Implementing custom sequencers is simpler since sequencers generate additional content by using the JCR API directly, rather than the proprietary graph API used in the previous release. Sequencer implementations are also able to register the node types programmatically, which simplifies the overall configuration for a repository.

## MIME Type Detection

There are two additional methods to the `org.modeshape.jcr.api.Binary` interface to obtain the MIME type. Your applications can use these methods to discover the MIME type for a binary value and,

for example, to set the "jcr:mimeType" property on the node.

The Service Provider Interface (SPI) for custom MIME type detectors has been removed. Instead, the Apache Tika framework is used, which has several MIME type detectors and provides its own SPI for custom detectors.

## Text Extractors

A simple Service Provider Interface (SPI) has been added for simple creation of custom extractors.

## Configuration

In this release, there is one JSON (previously XML) configuration file for each repository, each "deployed" to an engine. As a result, it is now possible to deploy and undeploy repositories dynamically even when the engine is running and other repositories are in use.

A JSON schema dictates the allowed structure of the configuration files. This means they can be validated:

```
Problems problems = config1.validate();
if ( problems.hasProblems() ) {
    // Output a summary of the problems (with line numbers) ...
    System.out.println(problems);
}
```

Sensible defaults are in place for settings that are not configured explicitly. The following is now a valid configuration:

```
{ "name" : "my-repo" }
```

You can now load configuration via the following means:

- read from a `java.io.File`
- read from a resolved `java.net.URL`
- read from a `String` containing a URL or a path to a file on the file system or classpath
- read a string containing the configuration

It is also possible to access the configuration of a running repository, change the configuration, and then update the running repository:

```
// Get the configuration ...
RepositoryConfiguration config = repository1.getConfiguration();

// Edit the configuration (which is a JSON document) to change a
value ...
Editor editor = config.edit();
editor.getOrCreateDocument(FieldName.STORAGE)
    .getOrCreateDocument(FieldName.BINARY_STORAGE);
    .setNumber(FieldName.MINIMUM_BINARY_SIZE_IN_BYTES,
newLargeValueSizeInBytes);
Changes changes = editor.getChanges();

// Apply the changes to the deployed repository ...
Future<Boolean> future = engine.update(config.getName(), changes);
```

```
// And optionally wait until the repository configuration is updated
...
future.get();
```



### Note

Many configuration changes can now be applied to a repository while it is running, but not everything. For example, changing *where* data is stored will apply only after the repository is restarted.

## Migrating Content

The database system now provides efficient backup and restore functionality that works at the repository level. Each backup contains all of the content in all workspaces for a single repository. Backups can be used to restore a repository to an earlier state, and it also serves as a mechanism for migrating from the previous release.

## Chapter 4. Backup and Restore

### 4.1. Migrating from a Previous Release

Backup and restore can be used to migrate content stored in the provided hierarchical database system.

This is the proposed way for users to migrate such data for release 6.



#### Note

When migrating from a version of JBoss Data Virtualization prior to version 6.0.0, a migration tool will be provided to backup the repository from the previous installation. The backup can be restored using the `restoreRepository` method on the new (and empty) repository.

### 4.2. The Repository Manager

The `org.modeshape.jcr.api.RepositoryManager` interface contains the backup and restore functions.

```
public interface RepositoryManager {  
  
    ...  
  
    Problems backupRepository( File backupDirectory ) throws  
    RepositoryException;  
  
    ...  
  
    Problems restoreRepository( File backupDirectory ) throws  
    RepositoryException;  
  
    ...  
  
}
```

The following code demonstrates how to access the repository manager from a standard authenticated JCR session:

```
javax.jcr.Repository repository = ...  
javax.jcr.Credentials credentials = ...  
String workspaceName = ...  
javax.jcr.Session session = repository.login(credentials,workspaceName);  
org.modeshape.jcr.api.Session msSession =  
(org.modeshape.jcr.api.Session)session;  
org.modeshape.jcr.api.RepositoryManager repoMgr =  
((org.modeshape.jcr.api.Session)session).getWorkspace().getRepositoryManager  
( );
```

**Note**

Which workspace is used by the session is not important.

### 4.3. Backup a Repository

The **backupRepository** method provided by the `org.modeshape.jcr.api.RepositoryManager` interface is used to create a backup of the entire repository, including all workspaces that existed when the backup was initiated.

This method blocks until the backup is completed, so it is the caller's responsibility to invoke the method asynchronously if that is desired.

When this method is called on a repository that is being actively used, all of the changes made while the backup process is underway will be included. At some point near the end of the backup process, however, additional changes will be excluded from the backup. This means that each backup contains a fully-consistent snapshot of the entire repository as it existed near the time at which the backup completed.

The following code demonstrates usage of the backup method:

```
org.modeshape.jcr.api.RepositoryManager repoMgr = ...
java.io.File backupDirectory = ...
Problems problems = repoMgr.backupRepository(backupDirectory);
if ( problems.hasProblems() ) {
    System.out.println("Problems restoring the repository:");
    // Report the problems (we'll just print them out) ...
    for ( Problem problem : problems ) {
        System.out.println(problem);
    }
} else {
    System.out.println("The backup was successful");
}
```

Each backup is stored on the file system in a directory that contains a series of GZIP-ed files (each containing representations of approximately 100K nodes) and a subdirectory in which all the large BINARY values are stored.

**Note**

It is the application's responsibility to initiate each backup operation. There currently is no way to configure a scheduled backup. Doing so would add significant complexity.

### 4.4. Restore a Repository

Once you have a complete backup on disk, you can then restore a repository back to the state captured within the backup using the **restoreRepository** method provided by the `org.modeshape.jcr.api.RepositoryManager` interface.

To do this, start a repository (or perhaps a new instance of a repository with a different configuration) and, before it is used by any applications, restore the content.

The following code demonstrates usage of the restore method:

```
org.modeshape.jcr.api.RepositoryManager repoMgr = ...
java.io.File backupDirectory = ...
Problems problems = repoMgr.restoreRepository(backupDirectory);
if ( problems.hasProblems() ) {
    System.out.println("Problems backing up the repository:");
    // Report the problems (we'll just print them out) ...
    for ( Problem problem : problems ) {
        System.out.println(problem);
    }
} else {
    System.out.println("The restoration was successful");
}
```

Once a restore succeeds, the newly-restored repository will be restarted and ready for use.



## Chapter 5. Supported Configurations

### 5.1. Supported Data Sources and Translators

The following table provides a list of data sources and translators that are supported by Red Hat.

**Table 5.1. Supported Data Sources and Translators**

Data Source	Driver	Translator
Actian Analytics Express 2.0	2.0	actian-vector
Amazon Redshift	postgresql 9.2	redshift
Apache Accumulo 1.5.0	N. A.	accumulo
Apache Cassandra 2.2.4	N. A.	cassandra
Apache Hive 2.0.0	2.0.0	hive
Apache Solr 4.9.0	N. A.	solr
Apache Spark 1.6.0	Hive 1.2.1	hive
Cloudera Hadoop 5.5.1	5.5.1	impala
EDS 5.x	current Teiid Driver	teiid
Files – delimited, fixed length	N. A.	file
Generic Datasource-JDBC ansi	N. A.	jdbc-ansi
Generic Datasource-JDBC simple (postgresql84)	postgresql 8.4	jdbc-simple
Greenplum 4.x	postgresql 9.0	greenplum
HBase 1.1	phoenix 4.5.1 HBase 1.1	hbase
Hortonworks Hadoop	2.3.4	hive
IBM DB2 10	4.12.55	db2
IBM DB2 9.7	4.12.55	db2
Informix 12.10	4.10.JC5DE	informix
Ingres 10	4.1.4	ingres
JBoss Data Grid 6.4 (library mode)	N. A.	infinispan-cache
JBoss Data Grid 6.4 (remote client - hotrod)	N. A.	infinispan-cache-dsl
LDAP/ActiveDirectory v3	N. A.	ldap
MariaDB	mysql 5.1.22	mysql5
ModeShape/JCR 3.1	3.8.4	modeshape
MongoDB 3.0	N. A.	mongodb
MS Access 2013	N. A.	access
MS Excel 2010	N. A.	excel
MS Excel 2013	N. A.	excel
MS SQL Server 2008	4.0.2206.100	sqlserver
MS SQL Server 2012	4.0.2206.100	sqlserver
MySQL 5.1	mysql 5.1.22	mysql5
MySQL 5.5	mysql 5.1.22	mysql5
Netezza 7.2.x	7.2.1.1	netezza
Oracle 11g RAC	12.1.0.2.0	oracle
Oracle 12c	12.1.0.2.0	oracle
PostgreSQL 8.4	postgresql 8.4	postgresql

Data Source	Driver	Translator
PostgreSQL 9.2	postgresql 9.2	postgresql
REST/JSON over HTTP	N. A.	ws
RHDS 9.0	N. A.	ldap
Salesforce.com API 22.0	NA	salesforce
Salesforce.com API 34.0	NA	salesforce-34
SAP_Hana 1.00.102.01.1444147999	1.00.82.00_03942 70-1510	hana
SAP Netweaver Gateway	odata 4	sap-nw-gateway
Sybase ASE 15	jconn4-26502	sybase
Sybase IQ 16 Express	jconn4-v7	sybaseIQ
Teradata Express 15	15.10.00.05	teradata
Vertica 7.2.1	7.2.1-0	vertica
Webservices	N. A.	ws
XML Files	N. A.	FILE



### Note

MS Excel is supported in so much as there is a write procedure.



### Note

The MySQL InnoDB storage engine is not suitable for use as an external materialization target.

## Appendix A. Revision History

<b>Revision 6.3.0-08</b> Updates for 6.3.	<b>Thu May 11 2017</b>	<b>David Le Sage</b>
<b>Revision 6.2.0-02</b> Updates for 6.2.	<b>Tue Aug 4 2015</b>	<b>David Le Sage</b>