



Red Hat JBoss Data Virtualization 6.4

Getting Started Guide

Learn how to perform a basic installation of Red Hat JBoss Data Virtualization and perform some rudimentary tasks with the product.

Red Hat JBoss Data Virtualization 6.4 Getting Started Guide

Learn how to perform a basic installation of Red Hat JBoss Data Virtualization and perform some rudimentary tasks with the product.

David Sage
dlesage@redhat.com

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

This guide teaches you how to perform a basic installation of the product.

Table of Contents

PREFACE	3
CHAPTER 1. GOALS OF THIS GUIDE	4
CHAPTER 2. KEY COMPONENTS	5
CHAPTER 3. MINIMUM HARDWARE REQUIREMENTS	6
CHAPTER 4. DOWNLOADING AND INSTALLING THE PRODUCT	7
4.1. PREREQUISITES	7
4.2. DOWNLOAD JBOSS DATA VIRTUALIZATION INSTALLER	7
4.3. INSTALL JBOSS DATA VIRTUALIZATION	7
CHAPTER 5. STARTING AND STOPPING RED HAT JBOSS DATA VIRTUALIZATION	10
5.1. STARTING RED HAT JBOSS DATA VIRTUALIZATION	10
5.2. STOPPING RED HAT JBOSS DATA VIRTUALIZATION	10
CHAPTER 6. FURTHER STEPS	11

PREFACE

Businesses increasingly need to access data residing in multiple disparate data sources. Therefore, you need to consider ways of making this information readily available for them to use.

Red Hat JBoss Data Virtualization is a data integration solution used to integrate data from sources such as relational databases, text files, web services, and ERP/CRM mainframe systems, as well as "big data" datasources such as Apache Hadoop (Hive) and MongoDB.

Red Hat JBoss Data Virtualization provides a unified virtualized view of information combined from multiple disparate sources. It hides details from the end user about the physical data sources, such as location, structure, API, access language and storage technology. This allows for more effort to be spent on data analysis and manipulation rather than on technical issues regarding the physical separation of the data. A virtual database (VDB) is used to map physical data sources to integrated views. At runtime, queries submitted against these views are coordinated among the dependent physical data sources, according to query criteria and the mappings defined by the VDB.

CHAPTER 1. GOALS OF THIS GUIDE

By following this guide, you will be taught how to download and install the product for testing in a non-production environment. (If you want to learn about installing on production systems, please refer to the *Red Hat JBoss Data Virtualization Installation Guide*.)

CHAPTER 2. KEY COMPONENTS

A complete JBoss Data Virtualization solution consists of the following:

The Server

The server is positioned between business applications and one or more data sources. It coordinates integration of these data sources so they can be accessed by the business applications at runtime.

Design Tools

Various design tools are available to assist users in setting up JBoss Data Virtualization for a particular data integration solution.

Administration Tools

Various management tools are available for administrators to configure and monitor JBoss Data Virtualization.

CHAPTER 3. MINIMUM HARDWARE REQUIREMENTS

- 16 GB Java Virtual Machine memory size
- A modern multi-core processor or a multi-socket system with multi-core processors
- 20 GB or more disk space for the JBoss server product and DV components. This will consist of:
 - 1 GB disk for installed product files
 - 5 GB or more for log files and deployed artifacts
 - 15 GB (default) for BufferManager maxBufferSpace
 - If Modeshape is used, you will need an additional 5 GB or more of filespace.

This will suffice for your development or testing system. For a more detailed discussion of system requirements when you later go to install Red Hat JBoss Data Virtualization on a production system, please refer to the Installation Guide.

CHAPTER 4. DOWNLOADING AND INSTALLING THE PRODUCT

4.1. PREREQUISITES

- Red Hat Customer Portal Account
- Adequate minimum system resources
- Administrative Security privileges for your chosen installation directory
- JAVA_HOME and PATH have been set in the Environment properties (for shortcuts to work on Microsoft Windows servers.)
- Java Developer Kit 7 or 8. (Note that if you intend to access Red Hat JBoss Data Grid 7.1, Red Hat JDV must be running JDK 1.8.)
- Java Virtual Machine. You can use the JVM installed with your JDK or an alternative such as OpenJDK.

4.2. DOWNLOAD JBOSS DATA VIRTUALIZATION INSTALLER

1. Go to <https://access.redhat.com/>
2. Click **Log in** and enter your **Red Hat Login** and **Password** to access the **Customer Portal**. You will need to register for an account if you do not yet have one.
3. Click **Downloads** ⇒ **Red Hat JBoss Data Virtualization**.
4. Click **Download** (next to the **Red Hat JBoss Data Virtualization [Version] Installer**) option.
5. Save the file.

4.3. INSTALL JBOSS DATA VIRTUALIZATION

Red Hat JBoss Data Virtualization is installed using a graphical wizard. The following instructions will guide you step-by-step through this graphical installer. In this example, we will be using default options to install the standalone version. The purpose of this Guide is to teach you how to install the product in a development or testing environment.

1. Open a command line window and navigate to the directory in which the GUI installer was downloaded.
2. Enter the following command to launch the GUI installer:

```
java -jar jboss-dv-VERSION-installer.jar
```

3. A dialogue box will open followed by the End User License Agreement. If you accept the terms of the agreement, click **I accept the terms of this license agreement** and then click **Next**.
4. Tell Red Hat JBoss Data Virtualization where Red Hat JBoss EAP is installed on your server or specify a new location if you do not have it installed as it comes bundled with the product. (If you have a pre-existing installation of Red Hat JBoss EAP, ensure that it is patched to the latest version of 6.4.x.) Click **Next**.

5. Ensure **Teiid Installation** and **Dashboard Builder** are selected. Click **Data Services Builder** if you want the graphical design tool as well or **Modeshape Installation** if you want that (you cannot have both Data Services Builder and Modeshape). Click **Next**.
6. You will be prompted to create a new EAP Admin, Dashboard Admin, Teiid data access user, and Modeshape user (if Modeshape is selected). You will also be asked if you want to enable OData access and a logging role. Once created, these are added to the AdminRealm and can be used to access the Management Console and other applications secured using ManagementRealm. Enter the new username and password in the appropriate fields and click **Next**.



IMPORTANT

Ensure that you remember all of these passwords. They give you access to different parts of the system. The EAP account and password allows you to administer the EAP Server, the Dashboard password is for administrative functions related to the web interface, and the Teiid data access user and ModeShape user are for standard user access.



NOTE

The username and password are not allowed to match and the password must have at least eight characters, with one alphabetical character, one numeric character, and one non-alpha-numeric character.

7. Select **Perform default configuration**. The database logging configuration is the only option selected by default. Click **Next**.
8. The **Configure password vault** screen appears. Input your desired password, which must have no fewer than six characters. Click **Next**.



NOTE

The H2 database (used as the default for the Vault) is not suitable for production databases. Use it in testing and evaluation environments only.

9. If you selected the Data Services Builder option, you will be prompted for SSL security information. Click "Generate new keystore." Leave the SSL Keystore alias as "jboss". Input a password of six characters. Leave the filepath as the default. Click **Next**.
10. A summary of the installation is displayed. Click **Next** for the installation to commence. This may take a few moments. Once all the components are installed, click **Next** again.
11. Click **Generate an installation script and properties file** if you wish to generate a script for automatic future installations. This script allows you to quickly reinstall or mass-deploy the product using the settings you have configured during the initial installation, without having to step through the wizard each time.
12. Make a note of the URLs of the Administration Console and the Data Virtualization Dashboard and click **Done** to exit the installer.

**NOTE**

Note that, after installing Red Hat JBoss Data Virtualization, if you move the product to another location, you may see some FileNotFound exceptions. This is because some filepaths are hard-coded by the JBoss EAP Server.

**WARNING**

If you attempt to use a vault with a keystore created with a different JDK than the one in which the data is stored, your server will fail to start. You must consistently use the same JDK when accessing the vault.

CHAPTER 5. STARTING AND STOPPING RED HAT JBOSS DATA VIRTUALIZATION

5.1. STARTING RED HAT JBOSS DATA VIRTUALIZATION

1. To run Red Hat JBoss Data Virtualization, you must first start the JBoss EAP server. To start the JBoss EAP server, follow the instructions below for your operating system:

- For Red Hat Enterprise Linux, open a command line window and enter this command from your EAP_HOME directory (the directory in which EAP is installed):

```
./bin/standalone.sh
```

- For Microsoft Windows, open a command line window and enter these commands:

```
chdir EAP_HOME/bin  
standalone.bat
```

2. To ensure that the server has started correctly and to verify that there have been no errors, check the server log: EAP_HOME/standalone/log/server.log



NOTE

You can also verify the execution was error-free by opening the Management Console in a web browser and logging in using the username and password of a registered JBoss EAP Management User. The console's address is <http://localhost:9990/console/> (For more information about using the Management Console see the *Red Hat JBoss Enterprise Application Platform Administration and Configuration Guide*.)



NOTE

For more advanced starting options, see the *Red Hat JBoss Enterprise Application Platform Administration and Configuration Guide*.

5.2. STOPPING RED HAT JBOSS DATA VIRTUALIZATION

To stop Red Hat JBoss Data Virtualization, halt the JBoss EAP Server. Do this by pressing Ctrl+C in the terminal in which EAP is running.

CHAPTER 6. FURTHER STEPS

Having worked through this book, Red Hat recommends that you read the Quick Starts Guide next as it provides information to help you start using this product. Quick starts are sample projects. By examining them, you can learn about functionality. You can also use them as templates for modeling your own business solutions.