



# **Red Hat JBoss Data Virtualization 6.2 Getting Started Guide**

---

David Sage



# Red Hat JBoss Data Virtualization 6.2 Getting Started Guide

---

David Sage  
dlesage@redhat.com

## Legal Notice

Copyright © 2017 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

---

## Table of Contents

<b>PREFACE</b> .....	<b>3</b>
<b>CHAPTER 1. INTRODUCTION AND GOALS OF THIS GUIDE</b> .....	<b>4</b>
<b>CHAPTER 2. A BROAD OVERVIEW OF THE PRODUCT</b> .....	<b>5</b>
<b>CHAPTER 3. MINIMUM HARDWARE REQUIREMENTS</b> .....	<b>6</b>
<b>CHAPTER 4. PRACTICAL TUTORIAL: DOWNLOADING AND INSTALLING THE PRODUCT</b> .....	<b>7</b>
4.1. PREREQUISITES .....	7
4.2. DOWNLOAD JBOSS DATA VIRTUALIZATION INSTALLER .....	8
4.3. INSTALL JBOSS DATA VIRTUALIZATION .....	8
<b>CHAPTER 5. PRACTICAL TUTORIAL: RUNNING RED HAT JBOSS DATA VIRTUALIZATION</b> .....	<b>11</b>
5.1. STARTING RED HAT JBOSS DATA VIRTUALIZATION .....	11
5.2. STOPPING RED HAT JBOSS DATA VIRTUALIZATION .....	11
<b>CHAPTER 6. CONCLUSION</b> .....	<b>12</b>



# PREFACE

## CHAPTER 1. INTRODUCTION AND GOALS OF THIS GUIDE

Welcome to **Red Hat JBoss Data Virtualization**. By following this guide, you will be taught how to download and install the product. First, though, you will be introduced to some basic concepts to teach you what Red Hat JBoss Data Virtualization does.

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

Data integration 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.

**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.



## CHAPTER 2. A BROAD OVERVIEW OF THE PRODUCT

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

To run a Red Hat JBoss Data Virtualization server, this is the very minimum hardware you will need:

1. 16 GB Java Virtual Machine memory size
2. A modern multi-core processor or a multi-socket system with multi-core processors
3. 20 GB or more disk space for the JBoss server product and DV components. This will consist of:
  - a. 1 GB disk for installed product files
  - b. 5 GB or more for log files and deployed artifacts
  - c. 15 GB (default) for BufferManager maxBufferSpace
  - d. If Modeshape is used, you will need an additional 5 GB or more of filesystem.

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. PRACTICAL TUTORIAL: DOWNLOADING AND INSTALLING THE PRODUCT

Now that you know a little about what the product can do, it is time to learn how to install it. First, you must make sure your server meets the pre-requisite requirements.

### 4.1. PREREQUISITES

Name	Description
Red Hat Customer Portal Account	Set up an account on the customer portal at <a href="https://access.redhat.com/">https://access.redhat.com/</a> . (It is important to ensure your Red Hat subscriptions are always kept up to date.)
System resources	Review the supported configurations table and ensure your system is supportable.
Security privileges	Ensure that you have administration privileges for your chosen installation directory.
Environment properties	Ensure that JAVA_HOME and PATH have been set in the Environment properties for shortcuts to work on Microsoft Windows servers.
Java Developer Kit	You must have Java 8 JDK installed on your system.
Red Hat JBoss Enterprise Application Platform	You must have Red Hat JBoss EAP 6.4 installed and patched to version 6.4.3. Download it from the Customer Portal and refer to the EAP documentation for installation instructions. Make sure you use the Installer version of EAP (not the ZIP version) and do a default installation. You can find Getting Started instructions for that product here: <a href="https://access.redhat.com/site/documentation/en-US/JBoss_Enterprise_Application_Platform/6.4/html/Getting_Started_Guide/index.html">https://access.redhat.com/site/documentation/en-US/JBoss_Enterprise_Application_Platform/6.4/html/Getting_Started_Guide/index.html</a>

Java Virtual Machine	Use the JVM installed with your JDK or use OpenJDK or the non-open source Azul Zing.
Red Hat JBoss Data Virtualization	You must have Red Hat Data Virtualization

## 4.2. DOWNLOAD JBOSS DATA VIRTUALIZATION INSTALLER

1. Login to the Red Hat Customer Portal.
  - a. Enter <https://access.redhat.com/> into the address bar of a browser.
  - b. 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.
2. Download JBoss Data Virtualization Installer
  - a. Click **Downloads** ⇒ **Red Hat JBoss Data Virtualization**.
  - b. Click **Download** (next to the **Red Hat JBoss Data Virtualization [Version] Installer**) option.
  - c. Save the file.

## 4.3. INSTALL JBOSS DATA VIRTUALIZATION

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

1. Open a terminal window and navigate to the location where the GUI installer was downloaded.
2. Enter the following command to launch the GUI installer:
 

```
java -jar jboss-dv-installer-VERSION.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. Add the filepath and click **Next**.
5. Ensure **Teiid Installation** and **Modeshape Installation** are selected. Click **Next**.
6. You will be prompted to create a new EAP Admin, Dashboard Admin, Teiid data access user and ModeShape user and whether you want to enable OData access. 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**.

You must 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 that 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-alphanumeric character.

7. You can install Red Hat JBoss Data Virtualization either with default configuration or with additional configuration options. For this exercise, we will be using the defaults only, so select **Perform default configuration** to install Red Hat JBoss Data Virtualization with default options (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 default H2 database is not suitable for production databases. Use it in testing and evaluation environments only.

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**.

9. Click **Generate an automatic installation script** if you wish to generate automatic script. This 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.
10. Click **Done** to exit the installer.

### 4.3.1. Result

Red Hat JBoss Data Virtualization is successfully installed and configured.



#### Note

Note that after installing Red Hat JBoss Data Virtualization, if you move the product to another location, you may see some FileNotFoundException exceptions. This is due the fact that 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.

If you want to learn about installation in more depth, please refer to the *Red Hat JBoss Data Virtualization Installation Guide*.

## CHAPTER 5. PRACTICAL TUTORIAL: RUNNING RED HAT JBOSS DATA VIRTUALIZATION

Now that you have installed the product, it is time to learn how to start and stop it.

### 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:

#### 5.1.1. Red Hat Enterprise Linux

1. Open an XTerm and enter this command from your EAP\_HOME directory:

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

#### 5.1.2. Microsoft Windows

1. Open a command-line window and enter these commands:

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

To ensure that the server has started correctly and to verify that there have been no errors, check the server log found at this location using a text editor: EAP\_HOME/MODE/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. CONCLUSION

By working through this guide, you have already learned quite a lot about the product. You have a basic overview of some key components and you now have a general understanding of the kind of problem Red Hat JBoss Data Virtualization can solve for your business. You also know how to perform a default installation and how to start and stop the software and check for errors.

Next, we recommend that you read the Quick Starts Guide 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.