



Red Hat JBoss Data Virtualization 6.4

Quick Starts Guide

Red Hat JBoss Data Virtualization 6.4 Quick Starts Guide

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

These quick starts give an overview of product functionality and can be adapted for your own needs.

Table of Contents

PREFACE	3
CHAPTER 1. INTRODUCTION AND GOALS OF THIS GUIDE	4
CHAPTER 2. WHAT ARE THE QUICK STARTS?	5
2.1. PREREQUISITES	5
2.2. GENERAL TIPS ON RUNNING QUICK STARTS	5
2.3. SIMPLE CLIENT QUICK START	5
2.4. DYNAMIC VDB DATA FEDERATION QUICK START	5
2.5. DYNAMIC VDB DATA ROLES QUICK START	7
2.6. WEB SERVICES AS A DATA SOURCE QUICK START	8
2.7. JBOSS DATA GRID RUNNING AS A DATA SOURCE QUICK START	8
2.8. HIBERNATE ON TOP OF TEIID QUICK START	8
CHAPTER 3. CONCLUSION	9

PREFACE

CHAPTER 1. INTRODUCTION AND GOALS OF THIS GUIDE

Congratulations! You now have quite a bit of experience with Red Hat JBoss Data Virtualization: you can perform a basic installation, run the software and use the Dashboard Builder to connect to data sources.

It is now time to showcase some more of the product's features that you can utilise to build your own solutions going forward. These features are demonstrated in the various "quick starts" that come with the product.

CHAPTER 2. WHAT ARE THE QUICK STARTS?

The quick starts are sample projects. The purpose of these quick start samples is to illustrate various JBoss Data Virtualization features. You can find quick starts in EAP_HOME/quickstarts directory.

2.1. PREREQUISITES

Name	Description
Java Developer Kit	You must have a Java 6 JDK or newer installed on your system.
Java Virtual Machine	Use the JVM installed with your JDK or use OpenJDK or the non-open source Azul Zing.
Maven	You must have Maven installed on your machine. For information about installing this, please refer to the Red Hat JBoss EAP documentation at https://access.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/6.4/html/development_guide/sect-install_maven_and_the_jboss_maven_repository .
Red Hat JBoss Data Virtualization	You must have Data Virtualization installed and running. Refer to "Red Hat JBoss Data Virtualization Getting Started Guide" for details on how to do this.

2.2. GENERAL TIPS ON RUNNING QUICK STARTS

When you want to run a quick start remember the following points:

1. Each quick start has different requirements. These are documented in their individual README.md files.
2. If the server is running in a mode other than standalone mode, you can run only some of the quick starts.

2.3. SIMPLE CLIENT QUICK START

The Simple Client quick start teaches you how to connect to Red Hat JBoss Data Virtualization using either the JDBC driver or a data source. You can use the Simple Client quick start to test your SQL queries against any JBoss Data Virtualization instance with an active virtual database. You must have an active virtual database before you run this quick start.

To run the quick start, go to the EAP_HOME/quickstarts/simpleclient directory and open the README.md file in a text editor and read the instructions it contains.

2.4. DYNAMIC VDB DATA FEDERATION QUICK START

The Dynamic VDB Data Federation quick start teaches you how Red Hat JBoss Data Virtualization can federate data from different sources into one virtual database.

The quick start uses three sample data sources to illustrate this concept, namely

1. Account Information (stored in an H2 database)
2. Market Data (stored in a comma-separated value [CSV] text file)
3. Other Personal Holdings (stored in a Microsoft Excel file).

The virtual database that is deployed is called Portfolio. It defines a model for each data source; these will be named Accounts, MarketData and PersonalValuations, respectively. The VDB is deployed to the runtime, thereby making it accessible to the user application via JDBC.

To run the quick start, go to the `EAP_HOME/quickstarts/dynamicvdb-datafederation` directory and open the `README.md` file in a text editor and read the instructions it contains.

When you run the quick start's test query, it joins data from the Accounts and MarketData models. Red Hat JBoss Data Virtualization then accesses both the relational and non-relational sources, calculates the portfolio values, and returns the results.

The `dynamicvdb-datafederation/src/vdb/portfolio-vdb.xml` file defines the resources that can be accessed by the client application.

In this model, the "vdb" element defines the virtual database that has a name of "Portfolio" with version "1". A "model" element represents a schema that is being integrated. This sample defines three sources. The "source" element inside the "model" element defines the "name" of the source (can be any name), the name of the translator (defines the type of the source like oracle, db2, mysql, h2, file or webservice) and the "connection-jndi-name" defines the source's JNDI name in the Red Hat JBoss EAP container.

Also note that inside the "model" elements, some "property" elements can be used to define how metadata is to be imported from the source. When you configure the quick start, the `portfolio-vdb.xml` is deployed to the JBoss EAP server. The `portfolio-vdb.xml.dodeploy` file (also referred to as a marker file, which is empty) is also deployed so that the Teiid Runtime is triggered to load the virtual database. The VDB's data sources are validated for availability. If they are available, the VDB status is set to "Active".

There are two virtual models defined in the dynamic VDB; Stocks and OtherHoldings. The Stocks virtual model contains 2 basic views. One, Stocks, to access the stock prices in the text data source and the second, StockPrices, that will join the Stocks with Products to federate the data from both data sources. The OtherHoldings contains a view that joins Accounts with OtherHoldings to show to whom the holdings belong.

This example uses three data sources, a relational database, a CSV text file and an Excel file. You can substitute any other relational database for H2, as long as you have a suitable JDBC driver. The schema file provided at `dynamicvdb-datafederation/src/teiidfiles/customer-schema.sql`, is specific to H2 but you can convert it for use with other databases.

In order to use a text file as the source, we need a data file which defines the data in a CSV format (comma separated values). Each data file contains column information for the table. The column information is typically defined on line 1, as header line in the file, and all the following lines contain the actual rows of data. Each line corresponds to single row. A portion of the sample file is shown below. The complete sample file is `dynamicvdb-datafederation/src/teiidfiles/data/marketdata-price.txt`.

```
SYMBOL, PRICE
IBM, 83.46
RHT, 11.84
```

BA, 44.58
ORCL, 17.37

The Excel file contains sheet1, which defines a header row and subsequent data rows, similar to a text file. This example demonstrates that Excel files can be used on different operating systems, not just on Microsoft Windows. You can use the provided data files or create your own data files. The sample file is `dynamicvdb-datafederation/src/teiidfiles/excelFiles/otherholdings.xls`

2.5. DYNAMIC VDB DATA ROLES QUICK START

The Dynamic VDB Data Roles quick start demonstrates how you can use data roles to control access to data. (This includes read-only and read-write access.)

In this quick start, the virtual database (`portfolio-vdb.xml`) has three data access roles:

1. `read-only` - this role provides read-only access (in other words, selects). This role is given to everybody who has a login credential. (Use the user called "user" with password "user".)
2. `read-write` - this role gives write access, in addition to read access (insert/update/delete). This access is given only to users with the "superuser" JAAS role. (Try with the user called "portfolio" with the password "portfolio".)
3. `prices` - this role shows how you can limit viewing access to certain columns.

To run the quick start, go to the `EAP_HOME/quickstarts/dynamicvdb-dataroles` directory and open the `README.md` file in a text editor and read the instructions it contains. Look at the `application-roles.properties` and `application-users.properties` in the `EAP_HOME/MODE/configuration/` directory, because you have to use the `./add-user.sh` (bat) script (as indicated in the `README`).

Look at the `portfolio-vdb.xml` file to see how the quick start is used:

```
<data-role name="ReadWrite">
  <description>Allow reads and writes to tables and
procedures</description>

  <permission>
    <resource-name>Accounts</resource-name>
    <allow-create>>false</allow-create>
    <allow-read>>true</allow-read>
    <allow-update>>true</allow-update>
  </permission>

  <permission>
    <resource-name>MarketData</resource-name>
    <allow-create>>false</allow-create>
    <allow-read>>true</allow-read>
    <allow-update>>true</allow-update>
  </permission>

  <mapped-role-name>supervisor</mapped-role-name>
</data-role>
```

To see how the users and roles were defined for JAAS, look the `application-roles.properties` and `application-users.properties` files. JAAS uses `application-users.properties` file to determine user credentials. The `application-roles.properties` file maps the username to the `mapped-role-name` element.

2.6. WEB SERVICES AS A DATA SOURCE QUICK START

The Web Services as a Data Source quick start demonstrates how to use the Web Service Translator to call a web service that returns an XML document. The quick start also demonstrates how to translate the XML document into a relational database.

To run the quick start, go to the `EAP_HOME/quickstarts/webservices-as-a-datasource` directory, open the `README.md` file in a text editor and read the instructions.

When you run the quick start, the `CustomerRESTWebSvc.war` is deployed. This is the web service data source that JBoss Data Virtualization accesses. The dynamic VDB, namely `webservice-vdb.xml`, creates the VIEW using the Data Definition Language (DDL). It does so in order to read and transform the XML document returned from the web service.

2.7. JBOSS DATA GRID RUNNING AS A DATA SOURCE QUICK START

The JBoss Data Grid Running as a Data Source example demonstrates how to access a Red Hat JBoss Data Grid cache using the `infinispan-cache-dsl` translator.

To run the quick start, go to the `EAP_HOME/quickstarts/jdg-remote-cache` directory and open the `README.md` file in a text editor and read the instructions.

The source model metadata is created dynamically based on the configured resource adapter. Notice that there is no DDL. This quick start uses the dynamically created metadata by default instead. However, you can define your own metadata and provide any necessary overrides.

The quick start shows you how to do these things:

1. Define the source model to the object cache by specifying the translator type of "infinispan-cache-dsl"
2. Set up the `Infinispan-dsl` resource adapter in order to communicate to the remote cache
3. Configure the JBoss Data Grid Cache as a data source, for which the translator is mapped to via the `connection-jndi-name`. It also demonstrates the ability to do reads and writes (Insert, Update and Delete) to the cache. Example queries are provided, but the actual content in the Data Manipulation Language (DML) may not match the content you loaded in the cache.

There are two other quick starts that allow you to explore Red Hat JBoss Data Grid integration: `jdq7.1-remote-cache` and `jdq7.1-remote-cache-materialization`.

2.8. HIBERNATE ON TOP OF TEIID QUICK START

When you utilize Hibernate, you will normally have a one-object-to-one-data-source mapping. By using Red Hat JBoss Data Virtualization as the data source, you can integrate the object and the source at the data layer rather than application layer. This makes it easier to join related information and then expose that information through Hibernate.

The Hibernate on Top of Teiid quick start demonstrates how hibernate can take advantage of multiple data sources through a single Java object by using the data federation capabilities of Red Hat JBoss Data Virtualization. This quick start extends the Portfolio VDB and creates a view that is mapped to a single relationally mapped object in hibernate.

To run the quick start, go to the `EAP_HOME/quickstarts/hibernate-on-top-of-teiid` directory and open the `README.md` file in a text editor and read the instructions it contains.

CHAPTER 3. CONCLUSION

During this hour, you have been taken on a tour of Red Hat JBoss Data Virtualization's various features. You have run various quick starts that demonstrate the product's functionality and teach you a little about how it works.

You can now go forward and modify these quick starts to build your own solutions to suit your business needs. Use the pre-packaged quick starts as templates to guide you and refer to the full documentation suite to learn more about the product's capabilities.