



Red Hat JBoss Data Virtualization 6.4

Data Services Builder Guide

Red Hat JBoss Data Virtualization 6.4 Data Services Builder 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

This guide is intended for developers who require a simple web-based graphical tool to create and manage data services.

Table of Contents

1. INTRODUCTION	2
2. LAUNCH THE DATA SERVICES BUILDER	2
3. CREATE A DATA SERVICE FROM A SOURCE AND TEST IT	3
3.1. Create a data source	3
3.2. Create a data service	3
3.3. Manage your connections	3
3.4. Testing the data service	4
3.5. Test your data service with SQL statements	4
4. MANAGE YOUR DATA SOURCE	4
4.1. Download a data source	4
4.2. Edit an existing data source	4
4.3. Reuse an existing data source	5
4.4. Advanced Troubleshooting for a Data Source	5
5. MANAGE YOUR DATA SERVICE	5
5.1. Create a data service using two source tables	5
5.2. Edit a data service consisting of two tables	6
5.3. Perform advanced edits on a data service	6
5.4. Copy a data service	6
5.5. Export a data service	6
5.6. Import a data service	7
6. EXIT THE DATA SERVICES BUILDER	7

1. INTRODUCTION



IMPORTANT

The Data Services Builder is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information on Red Hat Technology Preview features support scope, see <https://access.redhat.com/support/offerings/techpreview/>.

The Data Services Builder provides a simple web-based user interface designed to create, edit, test and manage your data services.

A data service is a single view defined within a **vdb.xml** file, created from the metadata from one or two JDBC data source tables.



NOTE

You must be running JDK 1.7 or greater if you wish to use the Data Services Builder. The Data Services Builder is not installed by default. You must select it during the installation process.

2. LAUNCH THE DATA SERVICES BUILDER

Before you start using the Data Services Builder, you will need to launch it and test your JDBC connection to ensure it can communicate with your server.

1. Start the server in standalone mode: `./standalone.sh`
2. Go to <https://localhost:8443/ds-builder>
3. Enter your username and password to log in. Note that the user must be assigned the roles **user** and **odata**.
4. Upon entering the Data Service Builder for the first time, you will see the **Data Service** page. Click the Preferences icon (a gear symbol) to set your JDBC credentials.
5. Click **Teiid Settings**.
6. Enter your username and password.
7. On the preferences page, click **Update**.
8. Click **Test JDBC Connection** to verify your credentials.
9. Click **Return** to go back to the **Data Service** page.

If you forget your username or password, you can reset them in Red Hat JBoss EAP.

**NOTE**

When you click the **Configure Workspace Location** link on the login page you can customize your workspace.

3. CREATE A DATA SERVICE FROM A SOURCE AND TEST IT

The basic usage of the Data Service Builder is to create a data source which is then consumed by a data service, against which you can subsequently run queries. Connections which are used by the data sources can be managed.

3.1. Create a data source

You can use the Data Services Builder to quickly and simply create a new data source. You can then use this data source to build a service.

1. On the **Data Sources** page, click the **New Data Source** button.
2. Enter a name for your data source. You can also add an optional description.
3. Click on your connection from the Connections List.
4. A translator corresponding to the selected connection is displayed. You can change the translator if you wish.
5. By default, **Filter Connection** is unchecked. This allows the data source to include all of the exposed schemas in the database or catalog. For this example, click on **Filter Connection** to filter and expose just a part of the schema.
6. Click Next.
7. Click on a schema to see the list of available tables in the right panel. You can further filter the exposed tables by entering a filter in the Filter box.
8. Click **Finish**.

3.2. Create a data service

Now you can create a data service to consume the Data Source you just created.

1. On the **Data Service** page, click on the **New Data Service** button.
2. Enter a Name for your data service. You can also optionally add a description.
3. Click the arrow next to your data source to expand the list of tables.
4. Click on a table.
5. Click **Finish**.

3.3. Manage your connections

You can add more JDBC connections if you need them.

1. On the **Data Service** page, go to the left-hand navigation menu and click on **Connections**.

2. To add a connection, click **New**.
3. Click **Existing Driver Type** or import a new one by clicking on **New Driver**.
4. Enter a name for your connection.
5. Enter a unique JNDI identifier for your connection.
6. Click **Next** through the wizard and enter specific properties for your connection.
7. Click **Next** and wait until the connection is deployed.
8. Click **Finish**.

You can also edit or delete connections.

3.4. Testing the data service

Having created a data service, you can now test that it works by submitting a query and seeing if a result is returned.

1. On the **Data Service** page, click on the **Test** button on a row of your table. A service is deployed and the Test page launches.
2. Click **Submit** to execute your query. (By default the OData tab is selected.)
3. You can customize and filter your query using row limit, columns, 'where' criteria and order-by options.

3.5. Test your data service with SQL statements

You may want to run a SQL query to test your data service instead. This allows you to script advanced queries and recycle statements.

1. On the **Test Data Service** page, click on the SQL tab.
2. A default query is generated, which you can edit.
3. Click **Submit** to run the query.

4. MANAGE YOUR DATA SOURCE

In the first example, you were shown how to create a new data source. However, you may have an existing one that you wish to reuse or edit, in order to save time and reduce the risk of errors.

4.1. Download a data source

1. On the **Data Sources** page, click the context menu (the three dots) on the particular data source and select **Download**.
2. Click **Save**.

4.2. Edit an existing data source

1. On the **Data Sources** page, click on the particular data source.

2. Click **Edit**.
3. You can now modify the name, connection, schema, filter condition and so on.
4. Click **Save**.

4.3. Reuse an existing data source

Reusing a data source allows you to save time and avoid errors. To do so, you can copy an existing source and save it under a new name.

1. On the **Data Sources** page, select a data source.
2. Click the context menu (the three dots) and select **Copy**.
3. Enter a unique name and click **Save**.
4. Verify the data source is added to Data Source page.

4.4. Advanced Troubleshooting for a Data Source

Sometimes you may want to check the DDL for a source when you are troubleshooting.

1. On the **Data Sources** page, click on the context menu (the three dots) on the particular data source and select **Display DDL**.
2. Inspect the Data Source Schema DDL.
3. Click the window control when you want to close the DDL display panel.

5. MANAGE YOUR DATA SERVICE

5.1. Create a data service using two source tables

You may wish to create service utilizing two different tables. To do so, you will need to join them in some way to create a link.

1. On the **Data Service** page, click **New**.
2. Enter unique name for the data service and optionally a description.
3. Select two tables from the **Data Sources Explorer**. (You can also select tables from other data sources.)
4. Click **Next**.
5. Click to select at least one column from each table.
6. Click on a Join Type.
7. Click **Next**.
8. Specify zero or more criteria by adding AND/OR conditions.
9. For each condition, select left and right columns and a condition operator.

10. Click **Finish**.
11. Verify the new data services is shown on the **Data Service** page.

5.2. Edit a data service consisting of two tables

You may wish to change your columns or join type. You can quickly change them using a graphical wizard.

1. On the **Data Service** page, click on the data service and click **Edit**.
2. Click **Next**.
3. On this page, you can change which columns are selected and the join type. When you have finished, click **Next**.
4. On this site, you can change or delete existing conditions or add a new one. When you have finished, click **Finish**.

5.3. Perform advanced edits on a data service

An expert editing mode is provided. Use it to perform advanced changes to your data service joins.

1. On the **Data Service** page, click on your data service and click **Edit**.
2. Click the **Expert** tab.
3. Edit data service.
4. Click **Finish**.

5.4. Copy a data service

To save time or to deploy a service across multiple environments, you many want to copy and reuse it.

1. On the **Data Service** page, click the context menu (the three dots) and select **Copy**.
2. Enter a unique name for the copied data service.
3. Click **Save**.

5.5. Export a data service

You may wish to export your data service out of the Data Services Builder into a collaborative version control environment where your colleagues can work on it. Here is how to export it to Git.

1. On the **Data Service** page, click the context menu (the three dots) and click **Git Export**.
2. Select your Git configuration.
3. Input the repository's details.
4. Click **Next**.
5. Enter your username and password credentials.
6. Click **Export**.

7. Click **Finish**.

5.6. Import a data service

You can import a data service from a repository such as Git back into the Data Services Builder.

1. On the **Data Service** page, click **Import**.
2. Select your Git configuration.
3. Input the repository's details.
4. Click **Next**.
5. Enter your username and password credentials.
6. Click **Import**.
7. Click **Finish**.

6. EXIT THE DATA SERVICES BUILDER

Once you have finished your session, you will want to exit the Data Services Builder.

1. Click the top-right menu.
2. Click Logout. You are returned to the login screen.