



Red Hat JBoss Data Virtualization 6.2

Administration and Configuration Guide

This guide is for administrators.

Red Hat JBoss Data Virtualization 6.2 Administration and Configuration Guide

This guide is for administrators.

Red Hat Customer Content Services

Legal Notice

Copyright © 2017 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). 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

Read this guide to learn how to administer Red Hat JBoss Data Virtualization.

Table of Contents

PART I. INTRODUCTION	5
CHAPTER 1. READ ME	6
1.1. BACK UP YOUR DATA	6
1.2. VARIABLE NAME: EAP_HOME	6
1.3. VARIABLE NAME: MODE	6
1.4. RED HAT DOCUMENTATION SITE	6
PART II. INTRODUCTION TO ADMINISTERING JBOSS DATA VIRTUALIZATION	7
CHAPTER 2. ADMINISTRATION TOOLS	8
2.1. ADMINISTRATION TOOLS OVERVIEW	8
2.2. MANAGEMENT CONSOLE	8
2.3. MANAGEMENT CLI	12
2.4. ADMIN SHELL	14
2.5. JBOSS OPERATIONS NETWORK	24
2.6. DASHBOARD BUILDER	25
PART III. USER MANAGEMENT	31
CHAPTER 3. USER ACCOUNTS	32
3.1. USER ACCOUNTS	32
3.2. DATA ROLES	32
3.3. ADDING A JBOSS EAP MANAGEMENT USER	32
3.4. ADDING A MODESHAPE USER	36
PART IV. CONFIGURING THE PRODUCT	37
CHAPTER 4. VIRTUAL DATABASES	38
4.1. VIRTUAL DATABASES AND JBOSS DATA VIRTUALIZATION	38
4.2. VIRTUAL DATABASE DEPLOYMENT	38
4.3. DEPLOY A VDB VIA FILE DEPLOYMENT	39
4.4. DEPLOY A VDB VIA MANAGEMENT CONSOLE	39
4.5. DEPLOY A VDB VIA CLI	40
4.6. DEPLOY A VDB VIA ADMIN SHELL	40
4.7. DEPLOY A VDB VIA ADMIN API	41
4.8. VDB DEPENDENCIES	41
4.9. DATA SOURCE DEPLOYMENT	42
4.10. MANAGING DEPLOYED VIRTUAL DATABASES USING MANAGEMENT CONSOLE	61
CHAPTER 5. VERSIONING	62
5.1. VIRTUAL DATABASE VERSIONING	62
5.2. SET THE VDB VERSION	62
5.3. VIRTUAL DATABASE CONNECTION TYPE	62
5.4. SET THE VDB CONNECTION TYPE VIA ADMIN API	62
CHAPTER 6. RESOURCE ADAPTERS	63
6.1. RESOURCE ADAPTERS IN JBOSS DATA VIRTUALIZATION	63
6.2. CONFIGURING RESOURCE ADAPTERS	63
6.3. EXAMPLE CONFIGURATION	64
6.4. RESOURCE ADAPTER PROPERTIES	64
6.5. CONFIGURING RESOURCE ADAPTERS USING CLI SCRIPTS	64
6.6. FILE ADAPTER PROPERTIES	64
6.7. GOOGLE SPREADSHEET RESOURCE ADAPTER PROPERTIES	65

6.8. JBOSS DATA GRID RESOURCE ADAPTER PROPERTIES	66
6.9. LDAP ADAPTER PROPERTIES	67
6.10. SALESFORCE ADAPTER PROPERTIES	67
6.11. WEB SERVICES ADAPTER PROPERTIES	68
6.12. MONGODB ADAPTER PROPERTIES	69
CHAPTER 7. CXF CONFIGURATION	70
7.1. CXF CONFIGURATION	70
7.2. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE	70
7.3. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: WS-SECURITY	71
7.4. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: LOGGING	72
7.5. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: TRANSPORT SETTINGS	73
7.6. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: SSL SUPPORT (HTTPS)	74
7.7. CONFIGURE CXF FOR A SALESFORCE DATA SOURCE	75
CHAPTER 8. JDBC DATA SOURCES	77
8.1. CONFIGURING JDBC DATA SOURCES	77
8.2. EXAMPLE CONFIGURATION	77
8.3. INSTALL A JDBC DRIVER WITH THE MANAGEMENT CLI	77
8.4. INSTALL A JDBC DRIVER WITH THE MANAGEMENT CONSOLE	77
8.5. CREATE A NON-XA DATASOURCE WITH THE MANAGEMENT INTERFACES	78
8.6. CREATE AN XA DATASOURCE WITH THE MANAGEMENT INTERFACES	80
8.7. DATASOURCE PARAMETERS	82
8.8. JDBC DRIVER DOWNLOAD LOCATIONS	88
CHAPTER 9. LOGGING	90
9.1. OVERVIEW OF LOGGING	90
9.2. DEFAULT LOG FILE LOCATIONS	90
9.3. JBOSS DATA VIRTUALIZATION LOG CATEGORIES	90
9.4. COMMAND LOGGING	91
9.5. AUDIT LOGGING	93
9.6. ENABLE AUDIT AND COMMAND LOGGING	93
9.7. COMMAND LOG OUTPUT	93
CHAPTER 10. CLUSTERING	95
10.1. CLUSTERING IN JBOSS DATA VIRTUALIZATION	95
10.2. ENABLE CLUSTERING IN JBOSS DATA VIRTUALIZATION	95
10.3. CONFIGURE LOAD BALANCING	95
CHAPTER 11. CACHING	96
11.1. RESULT SET CACHING	96
PART V. MONITORING AND PERFORMANCE	97
CHAPTER 12. MONITORING	98
12.1. MONITORING JBOSS DATA VIRTUALIZATION	98
12.2. QUERY/SESSION DETAILS	98
12.3. SESSION/QUERY METRICS	98
12.4. BUFFER MANAGER METRICS	99
12.5. CACHE METRICS	100
CHAPTER 13. PERFORMANCE TUNING	101
13.1. USAGE TIPS	101
13.2. MEMORY MANAGEMENT CONSIDERATIONS	103
13.3. SCALABILITY CONSIDERATIONS	105

13.4. DISK USAGE CONSIDERATIONS	106
13.5. THREADING CONSIDERATIONS	107
13.6. CACHING CONSIDERATIONS	107
13.7. TRANSPORT CONSIDERATIONS	108
13.8. LARGE OBJECTS (LOBS)	109
13.9. LOB CONSIDERATIONS	110
13.10. OTHER PERFORMANCE TUNING CONSIDERATIONS	110
PART VI. REFERENCE	112
CHAPTER 14. GENERAL CONFIGURATION	113
14.1. JBOSS DATA VIRTUALIZATION SETTINGS	113
14.2. VIEWING JBOSS DATA VIRTUALIZATION SETTINGS USING MANAGEMENT CLI	113
14.3. CHANGING JBOSS DATA VIRTUALIZATION SETTINGS USING MANAGEMENT CLI	114
14.4. MANAGING TRANSPORT AND SSL SETTINGS USING MANAGEMENT CLI	114
14.5. MANAGING TRANSLATOR SETTINGS USING MANAGEMENT CLI	114
14.6. TRANSPORT SECURITY AUTHENTICATION MODES	115
14.7. MANAGING CORE CONFIGURATION USING JBOSS MANAGEMENT CONSOLE	115
14.8. PORTS USED BY RED HAT JBOSS DATA VIRTUALIZATION	116
14.9. CHANGE THE DEFAULT JDBC PORT USING MANAGEMENT CONSOLE	116
14.10. SYSTEM PROPERTIES	116
14.11. TEIID MANAGEMENT CLI	121
CHAPTER 15. DIRECTORY STRUCTURE	123
15.1. DIRECTORY STRUCTURE	123
APPENDIX A. REVISION HISTORY	125

PART I. INTRODUCTION

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.

PART II. INTRODUCTION TO ADMINISTERING JBOSS DATA VIRTUALIZATION

CHAPTER 2. ADMINISTRATION TOOLS

2.1. ADMINISTRATION TOOLS OVERVIEW

The following tools are available for you to use to administer Red Hat JBoss Data Virtualization:

Management Console

Management Command Line Interface (CLI)

AdminShell

Admin API

JBoss Operations Network

2.2. MANAGEMENT CONSOLE

2.2.1. Management Console and JBoss Data Virtualization

Red Hat JBoss Data Virtualization provides a plug-in for the web-based Management Console. It provides a web interface that allows you to configure and monitor Data Virtualization services deployed upon a running JBoss EAP instance.




NOTE

For more information about general use of the Management Console, not specific to JBoss Data Virtualization, see the *Red Hat JBoss Enterprise Application Platform Administration and Configuration Guide*.

2.2.2. Log in to the Management Console

Prerequisites

- JBoss EAP 6 must be running.
 - You must have already created a user with permissions to access the Console.
1. Launch your web browser and go to this address: <http://localhost:9990/console/App.html>



NOTE

Port 9990 is predefined as the Management Console socket binding.

2. Enter your username and password to log in to the Management Console.

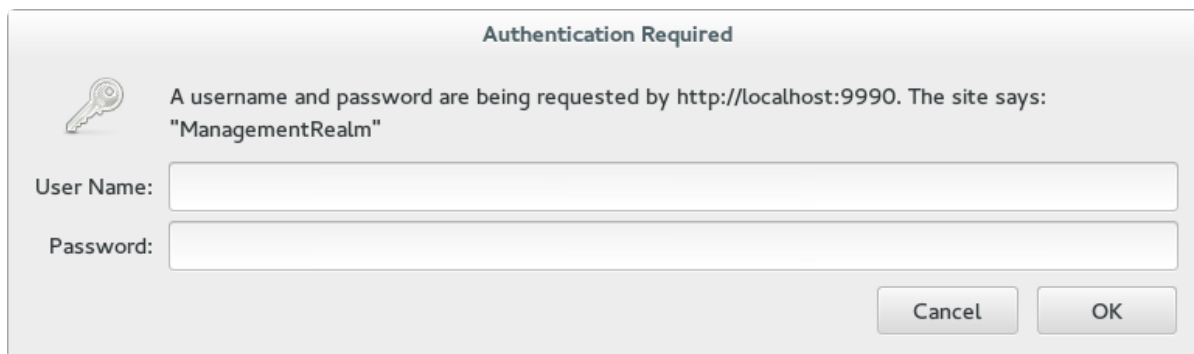


Figure 2.1. Log in screen for the Management Console

Result

Once logged in, you are redirected to the following address and the Management Console landing page appears: <http://localhost:9990/console/App.html#home>

2.2.3. Management Console - Configuration tab

The **Configuration** tab contains both general and Data Virtualization-specific configuration properties. Click **Teiid** in the left-hand navigational tree to view Teiid configuration. These are the sub-categories:

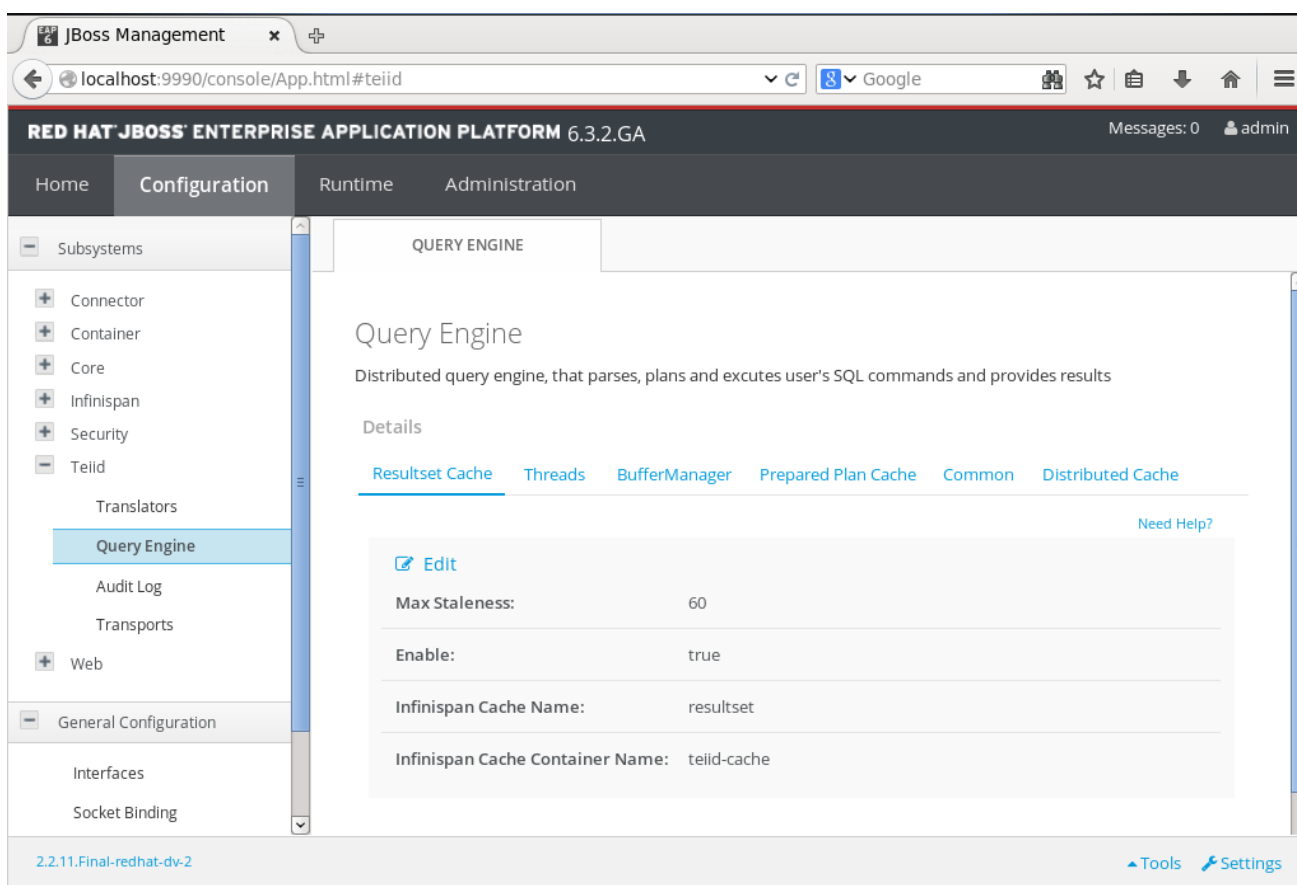


Figure 2.2. Configuration tab

- Query Engine - view and configure core Teiid engine properties.
- Translators - view and remove the Translators configured in Teiid.
- Transports - view, add and remove transports to the Teiid engine.

- Audit Logs - turn on or off audit logging for Teiid.

Use the **Configuration** tab to change any aspect of Teiid's configuration. Various different configuration properties are organised on different tabs. Click the **Need Help** link on each tab to see description of each field.



NOTE

Some properties require you to restart the server before changes you have made will take affect.

2.2.4. Management Console - Runtime tab

The **Runtime** tab shows runtime information about JBoss EAP server and Teiid subsystem. You can view runtime information about Teiid by clicking **Virtual Databases** in the left-hand side navigational tree.

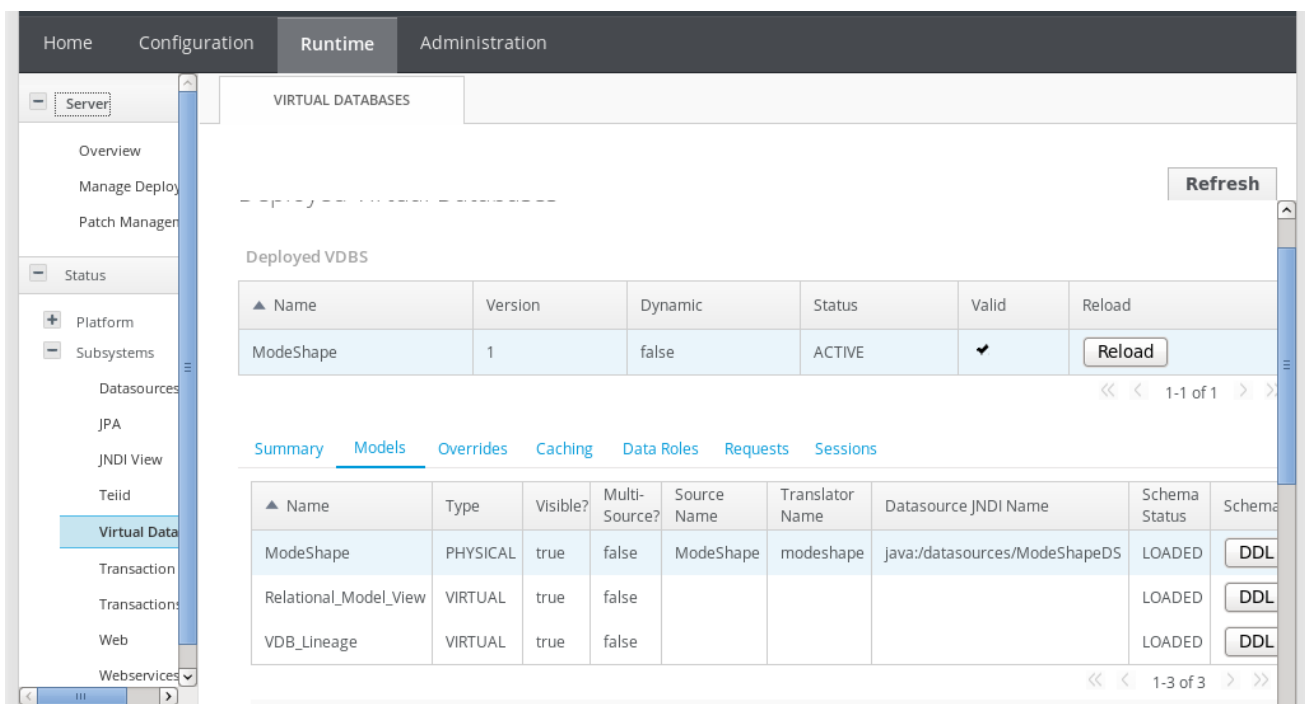


Figure 2.3. Runtime tab

All the VDBs deployed on the server are shown in top level table. When you select and highlight a VDB, more details about that VDB are displayed in the tabs below. Each of these tabs are divided into grouping of the functionality.

- **Summary:** This panel shows the details about the selected VDB and also any properties that are associated with that VDB, along with any other VDBs this VDB imports. This tab is designed to give a quick overview of the VDB status.
- **Models:** This panel shows all the models that are defined in a given VDB, and shows each model's translator name and source connection JNDI name. It also shows the type of models and if it is multi-source or not. When a particular model is selected it shows all properties of that model that are defined and also shows any errors associated with the model. When your VDB is not deployed in the active status, you must verify these errors and fix to resolve any deployment issues.

The **DDL** button in the Models panel shows the schema for the given model. Note that this does

not work for XML based models that are designed using Teiid Designer. The tool lets you edit the JNDI name by double clicking on them and modifying them. This useful if you would like to change the JNDI name in a given environment.

- **Overrides:** This panel shows the all the overridden translators in the Teiid Designer and their properties.
- **Caching:** Caching panel shows caching statistics of resultset cache as to how effectively the cache is being used. It also shows all the internal materialized views in the VDB and their load status as to when they were loaded. It also gives options to invalidate a specific view or all the views in a VDB, so that they can refresh or reload the contents from source. This panel also provides a UI to flush the entire the resultset cache contents or prepared plan cache contents for the selected VDB.
- **Data Roles:** Data Roles panel shows the all the policies that defined in the VDB using the Teiid Designer or hand coded in the vdb.xml file. For each selected policy, it also lists the permissions for that policy as to what kind of authorizations user has and shows the mapped enterprise role assignments to that policy. You can even add or remove an enterprise role from the policy using the this UI.
- **Requests:** This panel shows all the current requests against the selected VDB at the time of VDB selection. You can click refresh to get a more up to date requests. The top table in the panel shows the user submitted requests to the teiid engine, when one of those requests are selected, then the bottom table shows all the source level queries that are submitted to the physical sources by Teiid engine. Using this UI, user can also submit a cancel request to a user level query. Since cancel is an asynchronous operation, the operation is not guaranteed as query may already been finished, by the time cancel is submitted.
- **Sessions:** This panel shows all the active sessions that are connected to the selected VDB. It shows their connection properties and also gives an option to terminate either a selected session or all the sessions.

2.2.5. Management Console - Administration tab

You can configure roles and groups using Administration tab. Roles for a user to be included in and excluded from can be configured in the Management Console. You can also configure groups to be included or excluded from a role. Only users in the SuperUser or Administrator roles can perform this configuration.

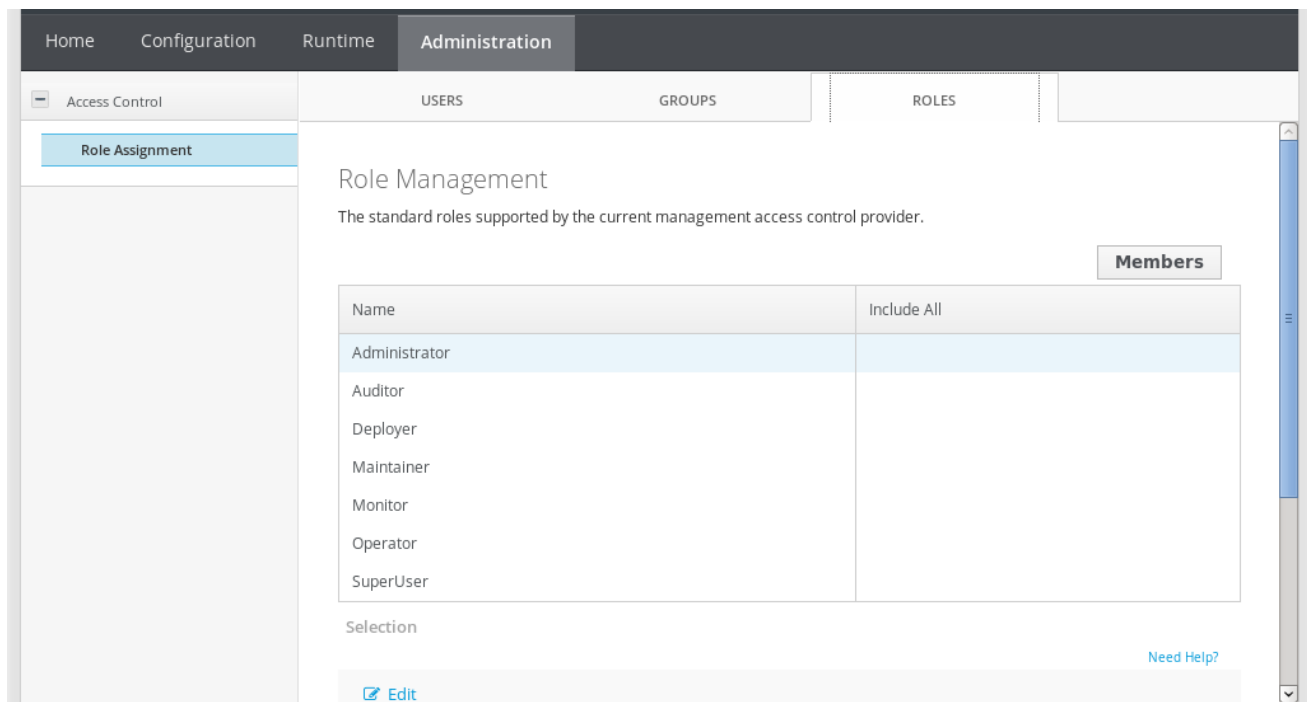


Figure 2.4. Administration tab

For more information related to User Management, see *Red JBoss EAP Administration and Configuration guide*.

2.3. MANAGEMENT CLI

2.3.1. Launch the Management CLI

Procedure 2.1. Launch CLI in Linux or Microsoft Windows Server

- ○ **Launch the CLI in Linux**
Run the `EAP_HOME/bin/jboss-cli.sh` file by entering the following at a command line:

```
$ EAP_HOME/bin/jboss-cli.sh
```

- **Launch the CLI in Microsoft Windows Server**
Run the `EAP_HOME\bin\jboss-cli.bat` file by double-clicking it, or by entering the following at a command line:

```
C:\>EAP_HOME\bin\jboss-cli.bat
```

2.3.2. Quit the Management CLI

From the Management CLI, enter the `quit` command:

```
[domain@localhost:9999 /] quit
```

2.3.3. Connect to a Managed Server Instance Using the Management CLI

Prerequisites

- [Section 2.3.1, “Launch the Management CLI”](#)

Procedure 2.2. Connect to a Managed Server Instance

- **Run the connect command**

From the Management CLI, enter the `connect` command:

```
[disconnected /] connect
Connected to domain controller at localhost:9999
```

- Alternatively, to connect to a managed server when starting the Management CLI on a Linux system, use the `--connect` parameter:

```
$ EAP_HOME/bin/jboss-cli.sh --connect
```

- The `--connect` parameter can be used to specify the host and port of the server. To connect to the address `192.168.0.1` with the port value `9999` the following would apply:

```
$ EAP_HOME/bin/jboss-cli.sh --connect --
controller=192.168.0.1:9999
```

2.3.4. Obtain Help with the Management CLI

Summary

Sometimes you might need guidance if you need to learn a CLI command or feel unsure about what to do. The Management CLI features a help dialog with general and context-sensitive options. (Note that the help commands dependent on the operation context require an established connection to either a standalone or domain controller. These commands will not appear in the listing unless the connection has been established.)

Prerequisites

- [Section 2.3.1, “Launch the Management CLI”](#)

1. For general help

From the Management CLI, enter the `help` command:

```
[standalone@localhost:9999 /] help
```

2. Obtain context-sensitive help

From the Management CLI, enter the `help -commands` extended command:

```
[standalone@localhost:9999 /] help --commands
```

3. For a more detailed description of a specific command, enter the command, followed by `--help`.

```
[standalone@localhost:9999 /] deploy --help
```

Result

The CLI help information is displayed.

2.4. ADMIN SHELL

2.4.1. AdminShell Features

AdminShell provides the following features:

Administration

AdminShell can be used to connect to a JBoss Data Virtualization instance in order to perform various administrative tasks.

Data Access

AdminShell can be used to connect to a virtual database (VDB) and run SQL commands to query VDB data and view results.

Migration

AdminShell can be used to develop scripts that will move VDBs and associated components from one development environment to another. (Users can test and automate migration scripts before executing them in production deployments.)

Testing

The built-in JUnit Test Framework allows users to write regression tests to check system health and data integrity. The written tests validate system functionality automatically, removing the need for manual verification by QA Personnel.

2.4.2. AdminShell Scripting

Use the Groovy language to write AdminShell scripts to automate tasks. AdminShell provides a custom library of Groovy functions to help administer JBoss Data Virtualization. These custom functions can be used in isolation or alongside other Groovy code.

Here are some basic rules. Memorise these to become familiar with the syntax:

- All commands and functions are case sensitive.
- Groovy commands are not required to end with a semicolon; this is optional.
- A function will take zero or more parameters as input.
 - Function parameters are provided within parentheses.
 - String parameters must be provided within single or double quotes.

```
connectAsAdmin("localhost", "9999", "user", "password", "conn1")
```

- Other Java classes and methods can be accessed and invoked from within scripts as long as required libraries are already in the class path. (JAR files added to the `lib` directory will be included in the class path automatically.)

```
import my.package.*;
myObject = new MyClass();
myObject.doSomething();
```



NOTE

You should disconnect from JBoss Data Virtualization before exiting AdminShell by entering the `disconnect()` command.

References

- For more information about Groovy scripting, refer to <http://groovy.codehaus.org/>.
- For more information about Groovy scripting and SQL, refer to <http://groovy.codehaus.org/Database+features>.
- For more information about testing using Groovy scripting, refer to <http://groovy.codehaus.org/Unit+Testing> and <http://junit.org/>.

2.4.3. AdminShell Help

All of the custom administrative methods available in AdminShell can be listed by using the `adminHelp()` method (note that only the custom AdminShell methods will be shown):

```
adminHelp()
```

If you require a specific method's definition and input parameters, use `adminHelp("METHOD")`:

```
adminHelp("deploy")

/*
 *Deploy a VDB from file
 */
void deploy(
String /* file name */)
throws AdminException
throws FileNotFoundException
```

Use the `sqlHelp()` method to list all SQL extension methods:

```
sqlHelp()
```

If you want to obtain a specific method's definition and input parameters, use `sqlHelp("METHOD")`.

2.4.4. AdminShell Basic Commands

Table 2.1. AdminShell Basic Commands

Command	Description
<code>adminHelp();</code>	Displays all of the custom AdminShell methods available.

Command	Description
<code>adminHelp("METHOD");</code>	Displays the definition and input parameters for the AdminShell method supplied.
<code>sqlHelp();</code>	Displays all of the custom SQL AdminShell methods available.
<code>sqlHelp("METHOD");</code>	Displays the definition and input parameters for the SQL AdminShell method supplied.
<code>println "STRING";</code>	Print something to console.
<code>sql = connect();</code>	Get an extended Groovy SQL connection using the default connection specified in the <code>connection.properties</code> file.
<code>sql.execute("SQL");</code>	Run any SQL command.
<code>connectAsAdmin();</code>	Connects as an admin user using the default connection specified in the <code>connection.properties</code> file. This command does not require a VDB name. Note that this is an admin connection, not a JDBC connection, so you cannot issue SQL commands using this connection. Note that if SSL is being used, you would need to adjust the connection URL and the client SSL settings as necessary (typically only for 2-way SSL).
<code>println getConnectionName();</code>	Returns the active connection name.
<code>useConnection("cNAME");</code>	Switches to using the given connection.
<code>disconnect();</code>	Disconnects the active connection.
<code>disconnectAll();</code>	Disconnects all connections (both active and suspended).

2.4.5. Execute a Script in Non-interactive AdminShell

Procedure 2.3. Execute a Script in Non-interactive AdminShell

1. Open a Command Line Terminal
2. Run the Script
Run the `/adminshell.sh load [-Dparam=value] PATH/FILENAME` command.



NOTE

Optional properties passed in using `-D` can be accessed from within the script file by using `System.getProperty()`.

```
value = System.getProperty("param")
```



IMPORTANT

Any connections established within the script should be disconnected before the script finishes.

2.4.6. Launch the Interactive AdminShell

Procedure 2.4. Launch the Interactive AdminShell

1. Navigate to the `adminshell` Directory

- a. Open a command line terminal.
- b. Navigate to the AdminShell directory: `cd /EAP_HOME/dataVirtualization/teiid-adminshell/`.

2. Unzip the AdminShell: `unzip teiid-VERSION-adminshell-dist.zip`

3. Launch the `adminshell` Script

Run the `./adminshell.sh` command.

The log output is as follows:

```

=====
==

Teiid AdminShell Bootstrap Environment

TEIID_HOME = /EAP_HOME/dataVirtualization/teiid-
adminshell-VERSION
CLASSPATH = /EAP_HOME/dataVirtualization/teiid-
adminshell-VERSION/lib/patches/*:/EAP_HOME/dataVirtualization/teiid-
adminshell-VERSION/lib/teiid-
adminshell-VERSION.jar:/EAP_HOME/dataVirtualization/teiid-
adminshell-VERSION/lib/*
JAVA = /usr/lib/jvm/java-1.7.0-openjdk.x86_64/bin/java

=====
==

==> [import static org.teiid.adminshell.AdminShell.*; import static
org.teiid.adminshell.GroovySqlExtensions.*; import
org.teiid.adminapi.*;]
Groovy Shell (1.7.2, JVM: 1.7.0_25)
Type 'help' or '\h' for help.
-----
-----
groovy:000>

```

Result

The interactive AdminShell is running. At this point you can execute individual commands line by line.

**NOTE**

You can exit the interactive AdminShell at any time by entering the `exit` command.

2.4.7. Helpful Tips for the Interactive AdminShell

Table 2.2. Interactive AdminShell Commands

Command	Description
↑ ('up arrow')	Displays previous commands executed, starting with the most recent.
<code>help</code>	Displays additional commands provided specifically for use within the interactive shell.

Shell Behavior

The interactive shell uses a special interpreter and displays behavior different from what would be expected from running a Groovy script:

- `def` statements do not define a variable in the context of the shell; For example, do not use `def x = 1`, use `x = 1`.
- The shell cannot parse Groovy classes that use annotations.

2.4.8. Save a Script in Interactive AdminShell

Procedure 2.5. Save a Script in Interactive AdminShell

1. **Open the Interactive AdminShell**
 - a. Open a command line terminal.
 - b. Navigate to `/EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/`.
 - c. Run the `./adminshell.sh` command.
2. **Start Recording to a File**
Enter the `record start PATH/FILENAME` command.
3. **Enter Desired Commands to Record**
Enter a series of commands for which you want to record the input/output.
4. **Stop Recording to the File**
Enter the `record stop` command.

Result

All input and output between the moments you issue the `record start` and `record stop` commands are captured in `PATH/FILENAME`.

**NOTE**

Since both input and output are recorded, the file will need some editing before it can be executed as a script file.

2.4.9. Execute a Script in Interactive AdminShell

Procedure 2.6. Execute a Script in Interactive AdminShell

1. Open the Interactive AdminShell

- a. Open a command line terminal.
- b. Navigate to `/EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/`.
- c. Run the `./adminshell.sh` command.

2. Run the Script

- o **Run the Script using the Interactive Shell load Command**
Run the `load PATH/FILENAME` command.
- o **Run the Script using the Groovy evaluate Command**
Run the `evaluate("PATH/FILENAME" as File)` command.

**NOTE**

You can also execute a script by entering it line by line after opening the interactive shell.

2.4.10. Launch the AdminShell GUI

Procedure 2.7. Launch the AdminShell GUI

1. Navigate to the adminshell Directory

- a. Open a command line terminal.
- b. Navigate to `/EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/`.

2. Launch the adminshell-console Script

Run the `./adminshell-console.sh` command.

Result

The AdminShell GUI is displayed.

**NOTE**

You can exit the AdminShell GUI at any time by clicking on **File** → **Exit**.

2.4.11. Save a Script in AdminShell GUI

Procedure 2.8. Save a Script in AdminShell GUI

Procedure 2.8. Save a Script in AdminShell GUI**1. Navigate to the adminshell Directory**

- a. Open a command line terminal.
- b. Navigate to `/EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/`.

2. Type a Script

Type your script in the upper *script window* of the AdminShell GUI. You may find it useful to test the script as you are preparing it by executing it via **Script** → **Run**.

3. Save the Script

- a. Select **File** → **Save As**.
- b. Choose a location and filename for the script and click on **Save**.

2.4.12. Execute a Script in AdminShell GUI

Procedure 2.9. Execute a Script in AdminShell GUI**1. Open the AdminShell GUI**

- a. Open a command line terminal.
- b. Navigate to `/EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/`.
- c. Run the `./adminshell-console.sh` command.

2. Prepare the Script**o Type a New Script**

Type your script in the upper *script window* of the AdminShell GUI.

o Load a Saved Script

- a. Select **File** → **Open**.
- b. Locate the desired script file and click on **Open**.

The script will appear in the upper *script window* of the AdminShell GUI.

3. Run the Script

Select **Script** → **Run**.

2.4.13. AdminShell Connection Properties

The `EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/connection.properties` file provides the default connection properties used by AdminShell to connect to an JBoss Data Virtualization instance:

```
jdbc.user=user
jdbc.password=user
jdbc.url=jdbc:teiid:admin@mm://localhost:31000;
admin.host=localhost
```



```
admin.port=9999
admin.user=admin
admin.password=admin
```

A call to `connect()` or `connectAsAdmin()` without any input parameters will connect to JBoss Data Virtualization using the settings defined in this properties file. `connect()` uses those properties prefixed with `jdbc` and `connectAsAdmin()` uses those properties prefixed with `admin`. Alternatively, you can include parameters in the `connect()` or `connectAsAdmin()` methods to connect using properties of your choice:

```
connect("URL", "USER", "PASSWORD")
```



WARNING

Do not leave passwords stored in clear text. The example above is for demonstration purposes only.

If you want to store your password in the file, take the necessary measures to secure it. Otherwise, do not use this feature at all and, instead, input your password interactively (or some other secure way.)

2.4.14. Multiple Connections in AdminShell

Using AdminShell, users can manage more than one connection to one or more JBoss Data Virtualization instances. For example, a user might want to have one connection to a development server and another to an integration server, simultaneously.

Every time a new connection is made, it is assigned a unique name and it becomes the active connection. If there is already a connection, it will be suspended (it will not be closed).

The `getConnectionName()` method returns the name of the active connection. The connection name can be assigned to a variable `cName` by the following command:

```
cName = getConnectionName();
```

The name of the current connection is required in order to switch from one to another. To change the active connection, use the `useConnection()` command, supplying the name (or a variable with the name assigned) of the connection you wish to use:

```
useConnection(cName);
```

Example

The following example demonstrates how to use two connections and switch between them:

```
// Create a connection
connectAsAdmin();

//Assign the connection name to conn1
```

```
conn1 = getConnectionName();

deploy("file.vdb")

// Create a second connection (which is now the active connection)
connectAsAdmin();

//Assign the new connection name to conn2
conn2 = getConnectionName();

deploy("file.vdb")

// Switch the active connection back to conn1
useConnection(conn1);

// Close the active connection (conn1)
disconnect();
```

2.4.15. Example Scripts

Example 2.1. Deploying a VDB

```
connectAsAdmin();
deploy("/path/to/<name>.vdb");

// check to validate the deployment
VDB vdb = getVDB("<name>", 1);
if (vdb != null){
    print (vdb.getName()+". "+vdb.getVersion()+" is deployed");
}
else {
    print ("<name>.vdb failed to deploy");
}
```

Example 2.2. Create a Datasource (Oracle)

```
connectAsAdmin();

// first deploy the JDBC jar file for Oracle
deploy("ojdbc6.jar");

props = new Properties();
props.setProperty("connection-url", "jdbc:oracle:thin:@<host>:1521:
<sid>");
props.setProperty("user-name", "scott");
props.setProperty("password", "tiger");

createDataSource("oracleDS", "ojdbc6.jar", props);
```

Example 2.3. Execute SQL Query against Teiid

```

sql = connect("jdbc:teiid:<vdb>@mm://<host>:31000", "user", "user");

// select
sql.eachRow("select * from sys.tables") { println "${it}" }

// update, insert, delete
sql.execute(<sql command>);

```

2.4.16. AdminShell FAQ

Q: Why won't the adminhelp command work in the GUI tool?

A: The AdminShell GUI environment does not understand shell commands such as **load**, **help**, and **adminhelp**, since they are not directly supported by Groovy. In the GUI you should use the equivalent functional form; for example, instead of using **adminhelp** use **adminHelp()**.

Q: Are there any pre-built scripts available?

A: Not currently.

Q: What is the difference between connectAsAdmin() and connect()?

A: The **connectAsAdmin()** method creates a contextual connection to the AdminAPI of JBoss Data Virtualization. The **connect()** method returns an extension of the Groovy SQL object to be used for SQL calls to JBoss Data Virtualization.

Q: What does getAdmin() do? Why do I need it?

A: The **getAdmin()** method returns the current contextual connection object that was created when you connected with **connectAsAdmin()**. This object implements the interface `org.teiid.adminapi.Admin`. The AdminShell commands provided are wrappers around this API. Advanced users can use this API directly if the provided wrapper commands do not meet their needs.

Q: Is IDE support available for writing the scripts?

A: The AdminShell GUI tool is a light-weight IDE. Full IDE support is available for Groovy, but requires manual manipulation of the class path and script imports.

Q: Can I use AdminShell methods in other environments?

A: The AdminShell methods (including the named contextual connections, AdminAPI wrapper, and help system) have no direct dependencies on Groovy and can be used in other scripting languages.

To use the AdminShell methods in another language, simply import the static methods and Admin classes into your script. You will also need to ensure that the `EAP_HOME/dataVirtualization/teiid-`

`adminshell-VERSION/lib/teiid-VERSION.jar` and `EAP_HOME/dataVirtualization/teiid-adminshell-VERSION/lib/teiid-adminshell-VERSION.jar` are in your class path.

The following snippet shows import statements that would work in Java, BeanShell, Groovy and others:

```
import static org.teiid.adminshell.AdminShell.*;
import org.teiid.adminapi.*;
```

Q: Is debugging support available?

A: The Interactive AdminShell and GUI have built-in support for inspection of the current state; however, line based debugging is beyond the scope of this document.

2.5. JBOSS OPERATIONS NETWORK

2.5.1. Installing JBoss Agent Plug-in Packs

JBoss Operations Network has additional agent plug-ins to manage other JBoss products. Although these are JBoss ON resource plug-ins, they are included in separate packages and require a separate subscription to download them.

1. Download the plug-in JAR files from the Customer Support Portal.
 - a. Navigate to <https://access.redhat.com/jbossnetwork/>.
 - b. Select JBoss ON for *PRODUCT*.
 - c. Select **Download**.
2. Extract the plug-in JAR files into the `JON_HOME/plugins` directory.
3. Have the JBoss ON server update its plug-ins. This can be done through the JBoss ON GUI or by restarting the server.
4. To load the plug-ins through the GUI:
 - a. Open the **Administration** tab.
 - b. In the **Configuration** area on the left, select the **Agent Plug-ins** link.
 - c. At the bottom of the list of loaded agent plug-ins, click the **SCAN FOR UPDATES** button.
5. Have the agents reload their plug-ins to load the new plug-ins. This can be done from the agent's command prompt using the `plugins` command:

```
> plugins update
```

**NOTE**

This can also be done in the JBoss ON GUI by scheduling an update plugins operation for an agent or a group of agents. Alternatively, restart the agents.

2.6. DASHBOARD BUILDER

2.6.1. Data sources

Red Hat JBoss Dashboard Builder can be connected to an external database, be it using JNDI of the container or connecting directly only using the JDBC driver to access the database. Connections to databases can be configured in workspace *Showcase* on page *External Connections*. After you have established the connection to the database, you need to create a data provider that will collect the data from the database and allow you to visualize it as an indicator in the dashboard area of a page.


When connecting to CSV files to acquire data, the connection is established directly through the data provider.

Note that Red Hat JBoss Dashboard Builder makes use of its own local internal database to store its local data. This database is read-only for Dashboard Builder, but is accessible from outside.

2.6.2. Connecting to data sources


You can connect either to a JNDI data source, that is, a data source set up and accessible from the application container, or directly to the data source as a custom data source, if the application container has the correct JDBC driver deployed.

To connect to an external data source, do the following:

1. Make sure the data source is up and running and that the application server has access to the data source. (Check the driver, the login credentials, etc. In Red Hat JBoss EAP 6, you can do so in the Management Console under **Subsystems** → **Connector** → **Datasources**)
2. In Dashboard Builder, on the Tree Menu (by default located on the of the Showcase perspective), go to **Administration** → **External connections**.
3. On the displayed External Connection panel, click the **New DataSource**  **Create new DataSource** button.
4. Select the data source type (JNDI or Custom DataSource) and provide the respective data source parameters below.

2.6.3. Creating data providers

To create a new data provider, do the following:

1. In the Tree Menu (the panel in the lateral menu of the Showcase workspace), click **Administration** → **Data providers**.
2. In the **Data Providers** panel, click the **Create new data provider**  **Create new data provider** button.

3. In the updated **Data Providers** panel, select in the **Type** dropdown menu the type of the data provider depending on the source you want the data provider to operate on.
4. Define the data provider parameters:

Data provider over a CSV file

- o Name: user-friendly name and its locale
- o CSV file URL: the url of the file (for example, `file:///home/me/example.csv`)
- o Data separator: the symbol used as separator in the CSV file (the default value is semicolon; if using comma as the separator sign, make sure to adapt the number format if applicable)
- o Quoting symbol: the symbol used for quotes (the default value is the double-quotes symbol; note that the symbol may vary depending on the locale)
- o Escaping symbol: the symbol used for escaping the following symbol in order to keep its literal value
- o Date format: date and time format
- o Number format: the format of numbers as resolved to thousands and decimals

Data provider over a database (SQL query)

- o Name: user-friendly name and its locale
- o Data source: the data source to query (the default value is `local`, which allows you to query the Dashboard Builder database)
- o Query: query that returns the required data

5. Click **Attempt data load**  to verify the parameters are correct.

6. Click **Save**.

7. In the table with the detected data, define the data type and if necessary provide a user-friendly name for the data. Click **Save**.

The data provider can now be visualized in an indicator on a page of your choice.

2.6.4. Dashboard Builder Workspace

2.6.4.1. Creating a workspace

To create a new workspace, do the following:

1. Click the **Create workspace** button on the top menu.



The management console with the **Workspace** node expanded and workspace management area with workspace details on the right is displayed.

2. In the **Create workspace** table on the right, set the workspace parameters:
 - Name: workspace name and its locale
 - Title: workspace title and its locale
 - Skin: skin to be applied on the workspace resources
 - Envelope: envelope to be applied on the workspace resources
3. Click **Create workspace**.
4. Optionally, click the workspace name in the tree menu on the left and in the area with workspace properties on the right define additional workspace parameters:
 - URL: the workspace URL
 - User home search: the home page setting

If set to **Role assigned page**, the home page configured in the page permissions is applied. Hence, you can set different home page for every role. (If you set it to **Current page**, all users will use the current home page as their home page.)

2.6.4.2. Creating Pages

To create a new page, do the following:


1. Make sure you are in the correct workspace.
2. Next to the **Page** dropdown box  in the top menu, click the **Create new page**  button.
3. The management console with the **Pages** node expanded and page management area with page details on the right is displayed.
4. In the **Create new page** table on the right, set the page parameters:
 - Name: page name and its locale
 - Parent page: parent page of the new page
 - Skin: skin to be applied on the page
 - Envelope: envelope to be applied on the page
 - Page layout: layout of the page
5. Click **Create new page**.
6. Optionally, click the page name in the tree menu on the left and in the area with workspace properties on the right define additional page parameters:
 - URL: the page URL
 - Visible page: visibility of the page

- Spacing between regions and panels

2.6.4.3. Defining Page permissions

Although users are usually authorized using the authorization method setup for the underlying application container (on Red Hat JBoss EAP, the **other** security domain by default), the Red Hat JBoss Dashboard Builder has its own role-based access control (RBAC) management tool to facilitate permission management on an individual page or multiple pages.

To define permissions on a page or all workspace pages for a role, do the following:

1. On the top menu, click the **General configuration**  button : the management console is displayed.
2. Under the **Workspace** node on the left, locate the page or the **Pages** node.
3. Under the page/pages node, click the **Page permissions** node.
4. In the **Page permissions** area on the right, delete previously defined permission definition if applicable and define the rights for the required role:
 - a. In the **Permission assignment** table, locate the **Select role** dropdown menu and pick the respective role.
 - b. In the **Actions** column of the table, enable or disable individual permissions.
5. Click **Save**.

2.6.4.4. Panels

A panel is a GUI widget, which can be placed on a page. There are three main types of panels:

Dashboard panels

are the primary BAM panels and include the following:

- **Data provider manager**: a panel with a list of available data providers and data provider management options
- **Filter and Drill-down**: a panel that displays all KPIs and their values to facilitate filtering in indicators on the given page defined over a data provider
- **HTML Editor panel**: a panel with static content
- **Key Performance Indicator (indicator)**: a panel that visualizes the data of a data provider

Navigation panels

are panels that provide navigation functions and include the following:

- **Breadcrumb**: a panel with the full page hierarchy pointing to the current page
- **Language menu**: a panel with available locales (by default in the top center)
- **Logout panel**: a panel with the name of the currently logged-in user and the logout button

- Page menu custom: a panel with vertically arranged links to all pages in the workspace (the list of pages can be adjusted) and general controls for the HTML source of the page
- Page menu vertical: a panel with vertically arranged links to all pages in the workspace (the list of pages can be adjusted)
- Page menu horizontal: a panel with horizontally arranged links to all pages in the workspace (the list of pages can be adjusted)
- Tree menu: a panel with the links to essential features such as Administration, Home (on the Home page of the Showcase workspace displayed on the left, in the lateral menu)
- Workspace menu custom: a panel with links to available workspaces (the list of workspaces can be adjusted) and general controls for the HTML source of the workspace
- Workspace menu horizontal: a horizontal panel with links to available workspaces (the list of workspaces can be adjusted)
- Workspace menu vertical: a vertical panel with links to available workspaces (the list of workspaces can be adjusted)


System panels

are panels that provide access to system setting and administration facilities and include the following:

- Data source manager: a panel for management of external data sources
- Export dashboards: a panel export of dashboards
- Export/Import workspaces: a panel for exporting and importing of workspaces

2.6.4.5. Adding panels

To add an existing panel to a page or to create a new panel, do the following:

1. Make sure the respective page is open (in the **Page** dropdown menu of the top menu select the page).
2. In the top menu, click the **Create a new panel in current page**  button.
3. In the displayed dialog box, expand the panel type you want to add (**Dashboard**, **Navigation**, or **System**) and click the panel you wish to add.
4. From the **Components** menu on the left, drag and drop the name of an existing panel instance or the **Create panel** item into the required location on the page.

If inserting a new indicator, the Panel view with the graph settings will appear. Define the graph details and close the dialog.

If adding an instance of an already existing indicator, you might not be able to use it, if it is linked to the KPIs on the particular original page. In such a case, create a new panel.

5. If applicable, edit the content of the newly added panel.

2.6.5. Dashboard Builder Filters

Filters work in the following way:

- You can define "shared" properties across several data set providers. A shared property is a property with the same id number in at least two different data providers.
- If you build a dashboard which refers to two or more data providers containing shared properties and try to filter by a shared property, then any of the KPIs containing this property will be filtered.
- Shared properties can be useful for implementing "join"-like filter behaviour. This allows you to filter several KPIs belonging to different data providers simultaneously.
- To disable the join-like behaviour, adjust the property ids on the **Data Provider Column Definition** screen. (Ensure the property ids are unique and do not clash between data provider definitions)

PART III. USER MANAGEMENT

CHAPTER 3. USER ACCOUNTS

3.1. USER ACCOUNTS

The following account types are available for you to use in JBoss Data Virtualization:

JBoss EAP Management User

You must have a JBoss EAP Management User in order to administer your JBoss Data Virtualization installation via the Management Console, Management CLI, JBoss Developer Studio, and Admin API.

JBoss Data Virtualization User

JBoss Data Virtualization users can access the virtual databases (VDBs). You can define what permissions each of these users has to dictate what they can do with these databases.

Hierarchical Database User

Hierarchical database users can access the provided hierarchical database. You can define what permissions each of these users has to dictate what they can do with this database.



NOTE

User/role names and details are independent of any other accounts such as operating system accounts. They relate only to JBoss EAP and JBoss Data Virtualization.

3.2. DATA ROLES

All authenticated users have access to a VDB. To restrict access, configure data roles. Set these in the Teiid Designer or the dynamic VDB's META-INF/vdb.xml file.

As part of the data role definition, you can map them to JAAS roles specified in `<mapped-role-name>` tags. (Establish these mappings using the `addDataRoleMapping()` method.)

How these JAAS roles are associated with users depends on which particular JAAS login module you use. For instance, the default `UsersRolesLoginModule` associates users with JAAS roles in plain text files.

For more information about data roles, see *Red Hat JBoss Data Virtualization Development Guide: Reference Material*.



IMPORTANT

Do not use "admin" or "user" as JAAS role names as these are reserved specifically for Dashboard Builder permissions.

3.3. ADDING A JBOSS EAP MANAGEMENT USER

3.3.1. Add the User for the Management Interfaces

The following procedure documents how to create the initial administrative user, in the event such a user is not created by the chosen installation method. This initial administrative user can use the web-

based Management Console and remote instances of the Management CLI to configure and administer JBoss EAP 6 from remote systems.

Procedure 3.1. Create the Initial Administrative User for the Remote Management Interfaces

1. Run the `add-user.sh` or `add-user.bat` script.

Change to the `EAP_HOME/bin/` directory. Invoke the appropriate script for your operating system.

Red Hat Enterprise Linux

```
[user@host bin]$ ./add-user.sh
```

Microsoft Windows Server

```
C:\bin> add-user.bat
```

2. Choose to add a Management user.

Press **ENTER** to select the default option **a** to add a Management user.

This user is added to the `ManagementRealm` and is authorized to perform management operations using the web-based Management Console or command-line based Management CLI. The other choice, **b**, adds a user to the `ApplicationRealm`, and provides no particular permissions. That realm is provided for use with applications.

3. Enter the desired username and password.

When prompted, enter the username and password. You will be prompted to confirm the password.

4. Enter group information.

Add the group or groups to which the user belongs. If the user belongs to multiple groups, enter a comma-separated list. Leave it blank if you do not want the user to belong to any groups.

5. Review the information and confirm.

You are prompted to confirm the information. If you are satisfied, type **yes**.

6. Choose whether the user represents a remote JBoss EAP 6 server instance.

Besides administrators, the other type of user which occasionally needs to be added to JBoss EAP 6 in the `ManagementRealm` is a user representing another instance of JBoss EAP 6, which must be able to authenticate to join a cluster as a member. The next prompt allows you to designate your added user for this purpose. If you select **yes**, you will be given a hashed secret value, representing the user's password, which would need to be added to a different configuration file. For the purposes of this task, answer **no** to this question.

7. Enter additional users.

You can enter additional users if desired, by repeating the procedure. You can also add them at any time on a running system. Instead of choosing the default security realm, you can add users to other realms to fine-tune their authorizations.

8. Create users non-interactively.

You can create users non-interactively, by passing in each parameter at the command line. This approach is not recommended on shared systems, because the passwords will be visible in log and history files. The syntax for the command, using the management realm, is:

```
[user@host bin]$ ./add-user.sh username password
```

To use the application realm, use the `-a` parameter.

```
[user@host bin]$ ./add-user.sh -a username password
```

9. You can suppress the normal output of the `add-user` script by passing the `--silent` parameter. This applies only if the minimum parameters `username` and `password` have been specified. Error messages will still be shown.

Result

Any users you add are activated within the security realms you have specified. Users active within the `ManagementRealm` realm are able to manage JBoss EAP 6 from remote systems.

3.3.2. Default User Security Configuration

Introduction

All management interfaces in JBoss EAP 6 are secured by default. This security takes two different forms:

- Local interfaces are secured by a SASL contract between local clients and the server they connect to. This security mechanism is based on the client's ability to access the local filesystem. This is because access to the local filesystem would allow the client to add a user or otherwise change the configuration to thwart other security mechanisms. This adheres to the principle that if physical access to the filesystem is achieved, other security mechanisms are superfluous. The mechanism happens in four steps:



NOTE

HTTP access is considered to be remote, even if you connect to the localhost using HTTP.

1. The client sends a message to the server which includes a request to authenticate with the local SASL mechanism.
 2. The server generates a one-time token, writes it to a unique file, and sends a message to the client with the full path of the file.
 3. The client reads the token from the file and sends it to the server, verifying that it has local access to the filesystem.
 4. The server verifies the token and then deletes the file.
- Remote clients, including local HTTP clients, use realm-based security. The default realm with the permissions to configure the JBoss EAP 6 instance remotely using the management interfaces is `ManagementRealm`. A script is provided which allows you to add users to this realm (or realms you create). For more information on adding users, see the *User Management* chapter of the *JBoss EAP 6 Administration and Configuration Guide*. For each user, the username and a hashed password are stored in a file.

Managed domain

```
EAP_HOME/domain/configuration/mgmt-users.properties
```

Standalone server

`EAP_HOME/standalone/configuration/mgmt-users.properties`

Even though the contents of the `mgmt-users.properties` are masked, the file must still be treated as a sensitive file. It is recommended that it be set to the file mode of `600`, which gives no access other than read and write access by the file owner.

3.3.3. Adding a JBoss Data Virtualization User

JBoss Data Virtualization's default configuration uses hashed files to define users and their roles. Follow these steps to add a new user for this default configuration.

1. Navigate to the bin director.
2. Launch the `./add-user.sh` script.
3. Selection option the "b) Application User (`application-users.properties`)" and add the username and password you desire.



IMPORTANT

The "admin" and "user" role names are reserved specifically for Dashboard Builder permissions.

3.3.4. Adding an Hierarchical Database User

JBoss Data Virtualization's default configuration uses plain text files to define users and their roles. Follow these steps to add a new user for this default configuration.



IMPORTANT

In previous versions of JBoss Data Virtualization, users could directly add a username and password in plain text to the `application-users.properties` file. The password is now hashed, and therefore adding this information directly into a file is no longer applicable or supported.

1. Use the `EAP_HOME/bin/add-user.sh` file to add a new username and password. For detailed steps, see [Section 3.3.1, "Add the User for the Management Interfaces"](#) and remember to select **b) Application User (`application-users.properties`)** instead of **a)** in the second step.
2. Save and close the file.
3. Open the `EAP_HOME/MODE/configuration/application-roles.properties` file in a text editor.
4. Add the user and assigned roles to the file as a new line with the syntax: `username=role1, role2, . . .`. Any number of roles can be assigned. If no roles are assigned to the user then no entry is required. For example:

```

administrator=administrator
harold=accounts, inventory, reports

```

5. Save and close the file.



NOTE

To temporarily disable user accounts, insert a hash (#) at the beginning of a line. When you remove the hash, functionality is restored.

3.4. ADDING A MODESHAPE USER

3.4.1. Create a Modeshape Publishing User

To use Modeshape publishing, you need to ensure that at least one user has the connect and readwrite roles or the connect and admin roles.

By default, all the roles are assigned to all anonymous-roles users.

PART IV. CONFIGURING THE PRODUCT

CHAPTER 4. VIRTUAL DATABASES

4.1. VIRTUAL DATABASES AND JBOSS DATA VIRTUALIZATION

JBoss Data Virtualization uses virtual databases (VDBs) to facilitate an enterprise data integration solution.

4.2. VIRTUAL DATABASE DEPLOYMENT

A virtual database (VDB) must be deployed before it can be accessed by client applications.

A VDB can be deployed by the following means:

File Deployment

File deployment is recommended for quick deployment during development, when the server is running locally on the developer's workstation.

Management Console

Deployment via the web-based Management Console is recommended for a simple way of deploying a VDB to a remote server.

Management Command Line Interface


Deployment via the EAP Management Command Line Interface (CLI) is another simple option for deployment.

AdminShell

Deployment via the AdminShell is recommended for more advanced deployments such as when you want to automate the deployment of artifacts in your environment.

Admin API

Deployment via the Admin API is recommended for more advanced deployments such as when you want to deploy a VDB from within other applications.



NOTE

A VDB can also be deployed within the Teiid Designer. See the *Red Hat JBoss Data Virtualization User Guide*.



IMPORTANT

Removal of a VDB cleans up all of its resources automatically; however, existing sessions will not be terminated automatically.

**WARNING**

When deploying a VDB, overwriting a VDB with one of the same name will terminate all connections to the old VDB unless VDB versioning is used.

Red Hat recommends the use of VDB versioning on production systems.

**WARNING**

You can locally name VDB artifacts as you wish, but the runtime names of your deployed VDB artifacts must have an extension of *.vdb for a zip file or *-vdb.xml for an xml file. Failure to name your file correctly will result in a deployment failure as the Teiid subsystem will not know how to handle it.

4.3. DEPLOY A VDB VIA FILE DEPLOYMENT

Prerequisites

- Red Hat JBoss Data Virtualization must be installed.

Procedure 4.1. Deploy a VDB via File Deployment

1. Copy your VDB into the `deploy` directory

Copy your VDB file into the `EAP_HOME/standalone/deployments` directory.

2. Create a marker file

Create an empty marker file of the same name with extension `.dodeploy` in the same directory. For example, if your VDB name is `enterprise.vdb`, then the marker file name must be `enterprise.vdb.dodeploy`.

**NOTE**

This only works in standalone mode. For domain mode, you must use one of the other available methods.

4.4. DEPLOY A VDB VIA MANAGEMENT CONSOLE

Prerequisites

- Red Hat JBoss Data Virtualization must be installed.
- The JBoss Enterprise Application Platform (EAP) server must be running.
- You must have a JBoss EAP Management User registered.

Procedure 4.2. Deploy a VDB via Management Console

1. **Launch the console in a Web browser**
Open <http://localhost:9990/console/> in a web browser.
2. **Authenticate to the console**
Enter your JBoss EAP administrator username and password when prompted.
3. **Open the Deployments panel**
In the **Runtime** view, select **Server** → **Manage Deployments**.
4. **Add the virtual database**
 - a. Select the **Add** button.
 - b. Select **Choose File** and choose the VDB file you want to deploy.
 - c. Select **Next** to review the deployment names then select **Save**.
 - d. Select **En/Disable** to enable the VDB.

4.5. DEPLOY A VDB VIA CLI

Prerequisites

- Red Hat JBoss Data Virtualization must be installed.
- The JBoss Enterprise Application Platform (EAP) server must be running.

Procedure 4.3. Deploy a VDB via CLI

1. **Open the command line interface**
Run the `EAP_HOME/bin/jboss-cli.sh` command.
2. **Connect to the server**
Run the `connect` command.
3. **Deploy the virtual database**
If in standalone mode, run `deploy PATH/DATABASE.vdb`.

If in domain mode, run `deploy -all-server-groups PATH/DATABASE.vdb`.



NOTE

In domain mode, you need to select a particular "server-group" or all available server groups are deployment options.

4.6. DEPLOY A VDB VIA ADMIN SHELL

Prerequisites

- Red Hat JBoss Data Virtualization must be installed.
- The JBoss Enterprise Application Platform (EAP) server must be running.

Procedure 4.4. Deploy a VDB via AdminShell

1. **Open the interactive AdminShell interface**
Run the `./adminshell.sh` command.
2. **Open a connection**
Within the interactive AdminShell, run the `connectAsAdmin()` command.
3. **Deploy your virtual database**
Run the `deploy("PATH/DATABASE.vdb")` command.
4. **Close the connection**
Run the `disconnect()` command.
5. **Exit the interactive shell**
 - a. Enter the `exit` command to leave the interactive shell.

**NOTE**

In domain mode, when deploying using AdminShell, the VDB is deployed to all available servers.

A VDB can also be deployed via the AdminShell console or using a script via the non-interactive AdminShell. For more information on using these, refer to topics on the Administration Shell.

4.7. DEPLOY A VDB VIA ADMIN API

You can deploy a VDB using the `deploy` method provided by the Admin interface within the Admin API package (`org.teiid.adminapi`).

Javadocs for Red Hat JBoss Data Virtualization can be found on the [Red Hat Customer Portal](#).

**NOTE**

In domain mode, when deploying using the Admin API, the VDB is deployed to all available servers.

4.8. VDB DEPENDENCIES

When deploying a virtual database (VDB) in JBoss Data Virtualization, you also have to provide dependency libraries and configuration settings for accessing the physical data sources used by your VDB. (You can identify all dependent physical data sources by looking in `META-INF/vdb.xml` within the `EAP_HOME/MODE/deployments/DATABASE.vdb` file.)

For example, if you are trying to integrate Oracle and file sources in your VDB, then you are responsible for providing both the JDBC driver for the Oracle source, and any necessary documents and configuration files that are needed by the file translator.

Data source instances may be shared between multiple VDBs and applications. Consider sharing connections to sources that are heavy-weight and resource-constrained.

Once you have deployed the VDB and its dependencies, client applications can connect using the **JDBC**

API. If there are any errors in the deployment, the connection attempt will fail and a message will be logged. Use the Management Console (or check the log files) to identify any errors and correct them so you can proceed. See *Red Hat JBoss Data Virtualization Development Guide: Server Development* for information on how to use JDBC to connect to your VDB.



WARNING

Some data source configuration files may contain passwords or other sensitive information. For instructions on how to avoid storing passwords in plaintext, refer to the *JBoss Enterprise Application Platform Security Guide*.

4.9. DATA SOURCE DEPLOYMENT

4.9.1. Accumulo Data Sources

Accumulo data sources use a Teiid specific JCA connector that is deployed into JBoss EAP during installation. There are many ways to create a Accumulo data source, using CLI, AdminShell, admin-console, etc. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

```
Execute the following command using the CLI once you connected to the Server. Ensure you provide the correct URL and user credentials. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB.
batch /subsystem=resource-adapters/resource-adapter=accumulo/connection-definitions=teiid:add(jndi-name=java:/accumulo-ds, class-name=org.teiid.resource.adapter.accumulo.AccumuloManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=accumulo/connection-definitions=teiid/config-properties=ZooKeeperServerList:add(value=localhost:2181) /subsystem=resource-adapters/resource-adapter=accumulo/connection-definitions=teiid/config-properties=Username:add(value=user) /subsystem=resource-adapters/resource-adapter=accumulo/connection-definitions=teiid/config-properties=Password:add(value=password) /subsystem=resource-adapters/resource-adapter=accumulo/connection-definitions=teiid/config-properties=InstanceName:add(value=instancename) /subsystem=resource-adapters/resource-adapter=accumulo/connection-definitions=teiid/config-properties=Roles:add(value=public) /subsystem=resource-adapters/resource-adapter=accumulo:activate runbatch
```

These are the properties defined in the RAR file:

Table 4.1. Properties

Property	Description	Required?	Default?
----------	-------------	-----------	----------

Property	Description	Required?	Default?
ZooKeeperServerList	A comma-separated list of zoo keeper server locations. Each location can contain an optional port, of the format host:port	True.	None.
ZooKeeperServerList	A comma-separated list of zoo keeper server locations. Each location can contain an optional port, of the format host:port	True.	None.
Username	Connection User's Name	True.	None.
Password	Connection User's password	True.	None.
InstanceName	Accumulo instance name	True.	None.
Password	Connection User's password	True.	None.
Roles	optional visibility for user, supply multiple with comma separated	False.	None.

To find out all the properties that are supported by this Accumulo Connector execute this command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=accumulo)`

4.9.2. Amazon SimpleDB Data Sources

SimpleDB data sources use a Teiid specific JCA connector that is deployed into EAP during installation. There are many ways to create a SimpleDB data source, using CLI, AdminShell, admin-console, and so forth. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute the following command using the CLI once you connected to the Server. Make sure you provide the correct access keys. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB. `batch /subsystem=resource-adapters/resource-adapter=simpledb/connection-definitions=simpledbDS:add(jndi-name=java:/simpledbDS, class-name=org.teiid.resource.adapter.simpledb.SimpleDBManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=simpledb/connection-definitions=simpledbDS/config-properties=AccessKey:add(value=xxx) /subsystem=resource-adapters/resource-`

```
adapter=simpledb/connection-definitions=simpledbDS/config-
properties=SecretAccessKey:add(value=xxx) /subsystem=resource-
adapters/resource-adapter=simpledb:activate runbatch
```

To find out all the properties that are supported by this SimpleDB Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=simpledb)`

4.9.3. Cassandra Data Sources



WARNING

Technology Preview features are not supported, may not be functionally complete, and are not intended for production use. These features are included to provide customers with early access to upcoming product innovations, enabling them to test functionality and provide feedback during the development process.

Cassandra's data sources use a Teiid specific JCA connector that is deployed into EAP during installation. There are many ways to create a Cassandra data source, using CLI, AdminShell, admin-console, etc. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute the following command using the CLI once you connected to the Server. Make sure you provide the correct URL and user credentials. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB: `batch /subsystem=resource-adapters/resource-adapter=cassandra/connection-definitions=cassandraDS:add(jndi-name=java:/cassandraDS, class-name=org.teiid.resource.adapter.cassandra.CassandraManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=cassandra/connection-definitions=cassandraDS/config-properties=Address:add(value=127.0.0.1) /subsystem=resource-adapters/resource-adapter=cassandra/connection-definitions=cassandraDS/config-properties=Keyspace:add(value=my-keyspace) /subsystem=resource-adapters/resource-adapter=cassandra:activate runbatch`

To find out all the properties that are supported by this Cassandra Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=cassandra)`

4.9.4. File Data Source

File data sources use a Teiid specific JCA connector that is deployed into EAP during installation. There are many ways to create the file data source, using CLI, AdminShell, admin-console, and so forth. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute following command using the CLI once you connected to the Server. Make sure you provide the correct directory name and other properties below. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB: `batch /subsystem=resource-adapters/resource-adapter=file/connection-definitions=fileDS:add(jndi-name=java:/fileDS, class-`


```
name=org.teiid.resource.adapter.file.FileManagedConnectionFactory,
enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-
adapter=file/connection-definitions=fileDS/config-
properties=Parentdirectory:add(value=/home/rareddy/testing/)
/subsystem=resource-adapters/resource-adapter=file/connection-
definitions=fileDS/config-properties=AllowParentPaths:add(value=true)
/subsystem=resource-adapters/resource-adapter=file:activate runbatch
```

To find out all the properties that are supported by this File Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=file)`

4.9.5. Google Spreadsheet Data Sources

You can authenticate your Google account in either of two ways: ClientLogin (which requires login and password) or OAuth (more complicated as it requires the user to supply a refresh token).

The Google JCA connector is named `teiid-connector-google.rar`. The examples include a sample `google.xml` file. The JCA connector has number of config-properties that facilitate authentication. The JCA connector connects to one spreadsheet:

Table 4.2. Configuration Properties

Property	Description
AuthMethod	Method to access Google. This property can be set to one of: ClientLogin, OAuth2. If the setting is ClientLogin it is necessary to provide additional two config properties Username, Password. If the setting is OAuth2 it is necessary to provide RefreshToken.
Username	Username to Google account. Applicable only if AuthMethod=ClientLogin
Password	Password to Google account. Applicable only if AuthMethod=ClientLogin
SpreadsheetName	Required property with name of the Spreadsheet that is datasource for this connector.
BatchSize	Maximum number of rows that can be fetched at a time. Defaults to 4096.

To obtain an OAuth Refresh Token, go to Google's authentication site:

```
https://accounts.google.com/o/oauth2/auth?
scope=https%3A%2F%2Fwww.googleapis.com%2Fauth%2Fdrive+https%3A%2F%2Fspread
sheets.google.com%2Ffeeds&redirect_uri=urn:ietf:wg:oauth:2.0:oob&response_
type=code&client_id=217138521084.apps.googleusercontent.com
```

Next, copy the authorization code into the following POST request and run it on command line. (The refresh token will be in the response.)

```
curl --data-urlencode code=<AUTH_CODE> \
--data-urlencode client_id=217138521084.apps.googleusercontent.com \
--data-urlencode client_secret=gXQ6-l0kEjE1lVcz7giB4Poy \
--data-urlencode redirect_uri=urn:ietf:wg:oauth:2.0:oob \
--data-urlencode grant_type=authorization_code
https://accounts.google.com/o/oauth2/token
```

4.9.6. Infinispan Data Sources

Infinispan data sources use a Teiid specific JCA connector that is deployed into EAP during installation. This connector can be configured to support the following modes of Infinispan caches:

Table 4.3. Registry Properties

Cache Type	How to Obtain This Cache
Local Cache	JNDI
Local Cache	Configuration file
Remote Cache	JNDI
Remote Cache	Specify one or more host porrts.
Local Cache	Specify one or more Hot Rod client property files.

You can configure these properties for this connector:

Table 4.4. Registry Properties

Property	Re qui re d	Property Template	Description
CacheTypeMap	Ye s	cacheName:className[;pkFieldName] [,cacheName:className[;pkFieldName]. .]	Map the root Java Object class name to the cache, and identify which attribute is the primary key to the cache.
module	No	Nil	Specify the JBoss AS module that contains the cache classes that were defined in CacheTypeMap
CacheJndiName	No	Nil	JNDI name to fine the CacheContainer

Property	Required	Property Template	Description
RemoteServerList	No	host:port[;host:port....]	Specify the host and ports that will be clustered together to access the caches defined in CacheTypeMap
ConfigurationFileNameForLocalCache	No	Nil	The Infinispan Configuration xml file for configuring a local cache.
HotRodClientPropertiesFile	No	Nil	The HotRod properties file for configuring a connection to a remote cache

There are many ways to create the data source, using CLI, AdminShell, admin-console and so forth. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes:

```
Execute following command using CLI once you connected to the Server. Make sure you provide the correct directory name and other properties below. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB:
batch /subsystem=resource-adapters/resource-adapter=infinispan/connection-definitions=infinispanDS:add(jndi-name=java:/infinispanDS, class-name=org.teiid.resource.adapter.infinispan.InfinispanManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=infinispan/connection-definitions=infinispanDS/config-properties=CacheTypeMap:add(value=trades:org.somewhere.Trade;tradeId) /subsystem=resource-adapters/resource-adapter=infinispan/connection-definitions=infinispanDS/config-properties=Module:add(value=org.somewhere) /subsystem=resource-adapters/resource-adapter=infinispan/connection-definitions=infinispanDS/config-properties=CacheJndiName:add(value=java:/myCache) runbatch
```

To find out all the properties that are supported by this Infinispan Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=infinispan)`



IMPORTANT

Avoid Classloading Issues: If you are using a servlet or other type of web application to create the DefaultCacheManager for the cache, be sure not to include the Infinispan jar dependencies in the application, but add their module dependencies instead.

4.9.7. Infinispan-DSL Data Sources

You must install the Red Hat JBoss Data Grid modules in order to use this translator. Note that this translator uses different JBoss Data Grid modules for remote and local modes. Please refer to the JBoss Data Grid documentation at https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Data_Grid/6.4/html/Infinispan_Query_Guide/chap-The_Infinispan_Query_DSL.html for more information about this. This JBoss Data Grid guide contains

some basic information on using the Infinispan Query DSL. The jar required for library mode is `infinispan-embedded-query`, and in remote client-server mode the jar is `infinispan-query`, which is installed by default with `jboss-datagrid-server`.

Infinispan-DSL data sources use a Teiid-specific JCA connector that is deployed into EAP during installation. This connector can be configured to support the following modes of Infinispan caches that will be accessed using the Hot Rod client:

Table 4.5. Properties

Cache Type	Property Name	How to Obtain Cache
Remote Cache	CacheJndiName	JNDI
Remote Cache	RemoteServerList	Server list: specify one or more host ports
Remote Cache	HotRodClientPropertiesFile	HotRod client properties file

This requires, at a minimum, JDG 6.3 or better, that is configured to use Google Protobuf for object serialization. This will enable Teiid to query the cache as if it was a local cache configured with Lucene indexes.

You can configure the following properties for this connector:

Table 4.6. Properties

Property	Required?	Property Template	Description
CacheTypeMap	Yes	cacheName:className[;pkFieldName] [,cacheName:className[;pkFieldName].]	Map the root Java Object class name to the cache, and identify which attribute is the primary key to the cache.
ProtobinFile	Yes	Nil	Path to the Google Protobin file that's packaged in a jar (such as <code>/quickstart/addressbook.protobin</code>)
MessageMarshallers	Yes	marshaller [,marshaller,..]	Contains Class names mapped its respective message marshaller, (class:marshaller,[class:marshaller,..]), that are to be registered for serialization
MessageDescriptor	Yes	Nil	Message descriptor class name for the root object in cache
module	No	Nil	Specify the JBoss EAP module that contains the cache classes that need to be loaded.

Property	Required?	Property Template	Description
CacheJndiName	No	Nil	JNDI name to find the CacheContainer
RemoteServerList	No	host:port[;host:port...]	Specify the host and ports that will be clustered together to access the caches.
HotRodClientPropertiesFile	No	Nil	The HotRod properties file for configuring a connection to a remote cache

Here is the XML code for a resource-adapter that is used to connect to the JDG remote-query quick start:

```
<resource-adapter id="infinispanRemQS">
  <module slot="main" id="org.jboss.teiid.resource-
adapter.infinispan.dsl"/>
  <connection-definitions>
    <connection-definition class-
name="org.teiid.resource.adapter.infinispan.dsl.InfinispanManagedConnectio
nFactory" jndi-name="java:/infinispanRemote" enabled="true" use-java-
context="true" pool-name="infinispanDS">
      <config-property name="CacheTypeMap">

addressbook:org.jboss.as.quickstarts.datagrid.hotrod.query.domain.Person;i
d

      </config-property>
      <config-property name="ProtobinFile">
        /quickstart/addressbook.protobin
      </config-property>
      <config-property name="MessageDescriptor">
        quickstart.Person
      </config-property>
      <config-property name="Module">
        com.client.quickstart.pojos
      </config-property>
      <config-property name="MessageMarshallers">
org.jboss.as.quickstarts.datagrid.hotrod.query.domain.Person:org.jboss.as.
quickstarts.datagrid.hotrod.query.marshallers.PersonMarshaller,org.jboss.a
s.quickstarts.datagrid.hotrod.query.domain.PhoneNumber:org.jboss.as.quicksta
rts.datagrid.hotrod.query.marshallers.PhoneNumberMarshaller,org.jboss.as
.quickstarts.datagrid.hotrod.query.domain.PhoneType:org.jboss.as.quickstar
ts.datagrid.hotrod.query.marshallers.PhoneTypeMarshaller
      </config-property>
      <config-property name="RemoteServerList">
        127.0.0.1:11322
      </config-property>
    </connection-definition>
  </connection-definitions>
</resource-adapter>
```

```

        </connection-definition>
    </connection-definitions>
</resource-adapter>

```

4.9.8. JDBC Data Sources

Here is an example highlighting configuring an Oracle data source. The process is nearly identical regardless of the database vendor. Typically the JDBC jar and the configuration like connection URL and user credentials change.

There are configuration templates for all the data sources in the `jboss-install/docs/teiid/datasources` directory. A complete description how a data source can be added into JBoss EAP is also described here. Here are two different ways to create a datasource:

Firstly, deploy the required JDBC jar file. For example, if you are trying to create an Oracle data source, first you need to deploy the "ojdbc6.jar" file: `deploy /path/to/ojdbc6.jar`



NOTE

If JBoss EAP is running in standalone mode, you can also manually copy this 'ojdbc6.jar' to the `jboss-install>/standalone/deployments` directory, to automatically deploy without using the CLI tool.

Now create a data source using this driver. There are many ways to create the datasource using CLI, AdminShell or admin-console. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

```

Execute the following command using the CLI once you connected to the Server. Make sure you
provide the correct URL and user credentials and edit the JNDI name to match the JNDI name you
used in VDB: /subsystem=datasources/data-source=oracle-ds:add(jndi-
name=java:/OracleDS, driver-name=ojdbc6.jar, connection-url=jdbc:oracle:thin:
{host}:1521:orcl,user-name={user}, password={password})
/subsystem=datasources/data-source=oracle-ds:enable

```

4.9.9. LDAP Data Sources

LDAP data sources use a Teiid specific JCA connector which is deployed into EAP during installation. There are many ways to create the ldap data source, using CLI, AdminShell, admin-console etc. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

```

Execute the following command using the CLI once you connected to the Server. Make sure you
provide the correct URL and user credentials. Add any additional properties required by the connector
by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI
name you used in VDB: batch /subsystem=resource-adapters/resource-
adapter=ldap/connection-definitions=ldapDS:add(jndi-name=java:/ldapDS, class-
name=org.teiid.resource.adapter.ldap.LDAPManagedConnectionFactory,
enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-
adapter=ldap/connection-definitions=ldapDS/config-
properties=LdapUrl:add(value=ldap://ldapServer:389) /subsystem=resource-
adapters/resource-adapter=ldap/connection-definitions=ldapDS/config-
properties=LdapAdminUserDN:add(value={cn=???,ou=???,dc=???)
/subsystem=resource-adapters/resource-adapter=ldap/connection-
definitions=ldapDS/config-properties=LdapAdminUserPassword:add(value={pass})

```

```
/subsystem=resource-adapters/resource-adapter=ldap/connection-
definitions=ldapDS/config-properties=LdapTxnTimeoutInMillis:add(value=-1)
/subsystem=resource-adapters/resource-adapter=ldap:activate runbatch
```

To find out all the properties that are supported by this LDAP Connector execute the following command in the CLI. `/subsystem=teiid:read-rar-description(rar-name=ldap)`

4.9.10. MongoDB Data Sources

MongoDB data sources use a Teiid-specific JCA connector that is deployed into EAP during installation. There are many ways to create a MongoDB data source, using CLI, AdminShell, admin-console and so forth. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute the following command using the CLI once you connected to the Server. Make sure you provide the correct URL and user credentials. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB: `batch /subsystem=resource-adapters/resource-adapter=mongodb/connection-definitions=mongodbDS:add(jndi-name="java:/mongoDS", class-name=org.teiid.resource.adapter.mongodb.MongoDBManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=mongodb/connection-definitions=mongodbDS/config-properties=RemoteServerList:add(value="{host}:27017") /subsystem=resource-adapters/resource-adapter=mongodb/connection-definitions=mongodbDS/config-properties=Database:add(value="{db-name}") /subsystem=resource-adapters/resource-adapter=mongodb:activate runbatch`

These are the properties that are defined in the RAR file:

Table 4.7. Properties

Property	Description	Required?	Default
RemoteServerList	A comma-separated list of server locations. Each location can contain an optional port, of the format host:port	Not applicable.	Not applicable.
Username	Connection User's Name	False	None.
Password	Connection User's password	False	None.
Database	MongoDB database name	True	None.

To find out all the properties that are supported by this MongoDB Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=mongodb)`

4.9.11. Phoenix Data Source

The following is an example for setting up Phoenix Data Sources, which is a precondition for Apache HBase Translator. In addition to the Data Sources set up, this article also covers mapping Phoenix tables to an existing HBase table and creating a new Phoenix table.

There are configuration templates for Phoenix data sources in the `EAP_HOME>/docs/teiid/datasources` directory. A complete description of how a data source can be added into Red Hat JBoss EAP is also described here.

Configuring a Phoenix data source is nearly identical to configuring JDBC Data Sources. The first step is deploying the Phoenix driver jar. Using the below CLI command to deploy Phoenix driver: `module add --name=org.apache.phoenix --resources=/path/to/phoenix-[version]-client.jar --dependencies=javax.api,sun.jdk,org.apache.log4j,javax.transaction.api /subsystem=datasources/jdbc-driver=phoenix:add(driver-name=phoenix,driver-module-name=org.apache.phoenix,driver-class-name=org.apache.phoenix.jdbc.PhoenixDriver`

The driver jar can be downloaded from the Phoenix document.

Next, create the Data Source based on the above deployed driver, which is also like creating a JDBC Data Source. Using the below CLI command to create the data source: `/subsystem=datasources/data-source=phoenixDS:add(jndi-name=java:/phoenixDS, driver-name=phoenix, connection-url=jdbc:phoenix:{zookeeper quorum server}, enabled=true, use-java-context=true, user-name={user}, password={password})`

Ensure the URL, Driver, and other properties are configured correctly: The JNDI name needs to match the JNDI name you used in the VDB. The driver name needs to match the driver you deployed in the above steps.

The URL needs to match the HBase zookeeper quorum server. Here is an example:

```
jdbc:phoenix [ :<zookeeper quorum> [ :<port number> ] [ :<root node> ] ],
'jdbc:phoenix:127.0.0.1:2181'
```

`user-name/password` - The user credentials for Phoenix Connection

Mapping a Phoenix table to an existing HBase table takes two steps. The first step is installing `phoenix-[version]-server.jar` to the classpath of every HBase region server. An easy way to do this is to copy it into the HBase lib. (For more details see the Phoenix documentation.)

Next, you must execute the DDL to map a Phoenix table to an existing HBase table. The DDL can either be executed via the Phoenix Command Line, or executed by JDBC.

The following is an example for mapping an existing HBase Customer with the following structure:

As depicted above, the HBase Customer table has 2 column families, `customer` and `sales`, and each has 2 column qualifiers, `name`, `city`, `product` and `amount` respectively. Map this table to Phoenix via DDL:

```
CREATE TABLE IF NOT EXISTS "Customer" ("ROW_ID" VARCHAR PRIMARY KEY,
"customer"."city" VARCHAR, "customer"."name" VARCHAR, "sales"."amount"
VARCHAR, "sales"."product" VARCHAR)
```

For more about mapping Phoenix tables to an existing HBase table please refer to the Phoenix documentation.

Creating a new Phoenix table is just like mapping to an existing HBase table. Phoenix will create any metadata (table, column families) that do not exist. Similar to the above example the DDL to create the Phoenix/HBase Customer table would be:

```
CREATE TABLE IF NOT EXISTS "Customer"("ROW_ID" VARCHAR PRIMARY KEY,
"customer"."city" VARCHAR, "customer"."name" VARCHAR, "sales"."amount"
VARCHAR, "sales"."product" VARCHAR)
```

4.9.12. Salesforce Data Sources

Salesforce data sources use a Teiid specific JCA connector that is deployed into EAP during installation. There are many ways to create the salesforce data source, using CLI, AdminShell, admin-console etc. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute the following command using the CLI once you connected to the server. Make sure you provide the correct URL and user credentials. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB: `batch /subsystem=resource-adapters/resource-adapter=salesforce/connection-definitions=sfDS:add(jndi-name=java:/sfDS, class-name=org.teiid.resource.adapter.salesforce.SalesForceManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=salesforce/connection-definitions=sfDS/config-properties=URL:add(value=https://www.salesforce.com/services/Soap/u/22.0) /subsystem=resource-adapters/resource-adapter=salesforce/connection-definitions=sfDS/config-properties=username:add(value={user}) /subsystem=resource-adapters/resource-adapter=salesforce/connection-definitions=sfDS/config-properties=password:add(value={password}) /subsystem=resource-adapters/resource-adapter=salesforce:activate runbatch`

To find out all the properties that are supported by this Salesforce Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=salesforce)`

4.9.13. Solr Data Sources

Solr data sources use a Teiid specific JCA connector that is deployed into EAP during installation. There are many ways to create a Solr data source, using CLI, AdminShell, admin-console, etc. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute the following command using the CLI once you connected to the Server. Make sure you provide the correct URL and user credentials. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB: `batch /subsystem=resource-adapters/resource-adapter=solr/connection-definitions=solrDS:add(jndi-name=java:/solrDS, class-name=org.teiid.resource.adapter.solr.SolrManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=solr/connection-definitions=solrDS/config-properties=url:add(value=http://localhost:8983/solr/) /subsystem=resource-adapters/resource-adapter=solr/connection-definitions=solrDS/config-properties=CoreName:add(value=collection1) /subsystem=resource-adapters/resource-adapter=solr:activate runbatch`

To find out all the properties that are supported by this Solr Connector execute the following command in the CLI: `/subsystem=teiid:read-rar-description(rar-name=solr)`

4.9.14. Web Service Data Sources

Web service data sources use a Teiid specific JCA connector that is deployed into EAP during installation. There are many ways to create the file data source, using CLI, AdminShell, admin-console etc. The example shown below uses the CLI tool, as this works in both Standalone and Domain modes.

Execute following command using the CLI once you connected to the Server. Make sure you provide the correct endpoint and other properties below. Add any additional properties required by the connector by duplicating the "connection-definitions" command below. Edit the JNDI name to match the JNDI name you used in VDB: `batch /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS:add(jndi-name=java:/wsDS, class-name=org.teiid.resource.adapter.ws.WSManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=EndPoint:add(value={end_point}) /subsystem=resource-adapters/resource-adapter=webservice:activate runbatch`

To find out all the properties that are supported by this Web Service Connector execute the following command in the CLI. `/subsystem=teiid:read-rar-description(rar-name=webservice)`. The Web Service Data Source supports specifying a WSDL using the Wsdl property. If the Wsdl property is set, then the ServiceName, EndPointName, and NamespaceUri properties should also be set. The Wsdl property may be a URL or file location or the WSDL to use. You must restart the server.

Table 4.8. Registry Properties

Property	Application	Required?	Default?	Description
EndPoint	HTTP and SOAP.	True	Not applicable.	URL for HTTP and service endpoint for SOAP.
SecurityType	HTTP and SOAP.	false	none	Type of Authentication to used with the web service. Allowed values are "None","HTTPBasic","WSecurity" and "Kerberos".
AuthUserName	HTTP and SOAP.	false	Not applicable.	Name value for authentication, used in HTTPBasic and WsSecurity.
AuthPassword	HTTP and SOAP.	false	Not applicable.	Password value for authentication, used in HTTPBasic and WsSecurity
ConfigFile	HTTP and SOAP.	False	Not applicable.	CXF client configuration File or URL.
EndPointName	HTTP and SOAP.	False	Teiid	Local part of the end point QName to use with this connection, needs to match one defined in CXF file.
ServiceName	SOAP	False	Not applicable.	Local part of the service QName to use with this connection.

Property	Application	Required?	Default?	Description
NamespaceUri	SOAP.	False	http://teiid.org	Namespace URI of the service QName to use with this connection
RequestTimeout	HTTP and SOAP.	False	Not applicable.	Timeout for request.
ConnectTimeout	HTTP and SOAP.	False	Not applicable.	Timeout for connection.
WSDL	SOAP.	False	Not applicable.	WSDL file or URL for the web service.

Each web service data source may choose a particular CXF config file and port configuration. The `ConfigFile` config property specifies the Spring XML configuration file for the CXF Bus and port configuration to be used by connections. If no config file is specified then the system default configuration will be used.

Only one port configuration can be used by this data source. You may explicitly set the local name of the port QName to use via the `ConfigName` property. The namespace URI for the QName in your config file should match your WSDL/namespace setting on the data source or use the default of `http://teiid.org`. See the CXF Documentation and the sections below on WS-Security, Logging, etc. for examples of using the CXF configuration file.

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:http-conf="http://cxf.apache.org/transport/http/configuration"
  xsi:schemaLocation="http://cxf.apache.org/transport/http/configuration
    http://cxf.apache.org/schemas/configuration/http-conf.xsd
    http://www.springframework.org/schema/beans
  >
  <http-conf:conduit name="{http://teiid.org}configName.http-conduit">
    <http-conf:client ConnectionTimeout="120000"
    ReceiveTimeout="240000"/>
  </http-conf:conduit>
</beans>
```

In the conduit name `{http://teiid.org}configName.http-conduit`, the namespace, `{http://teiid.org}`, may be set via the namespace `datasource` property. Typically that will only need done when also supplying the `wsdl` setting. The local name is followed by `.http-conduit`. It will be based upon the `configName` setting, with a default value of `teiid`.



NOTE

You do not have to use the Spring configuration to set timeouts. The `ConnectionTimeout` and `ReceiveTimeout` can be set via the resource adapter `connectTimeout` and `requestTimeout` properties respectively.

To enable the use of WS-Security, the SecurityType should be set to WSSecurity. At this time Teiid does not expect a WSDL to describe the service being used. Thus a Spring XML configuration file is not only required, it must instead contain all of the relevant policy configuration. And just as with the general configuration, each data source is limited to specifying only a single port configuration to use:

```
batch /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS:add(jndi-name=java:/wsDS, class-name=org.teiid.resource.adapter.ws.WSManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigFile:add(value=${jboss.server.home.dir}/standalone/configuration/xxx-jbossws-cxf.xml) /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigName:add(value=port_x) /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=SecurityType:add(value=WSSecurity) /subsystem=resource-adapters/resource-adapter=webservice:activate runbatch
```

This is the corresponding `xxx-jbossws-cxf.xml` file that adds a timestamp to the SOAP header

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://cxf.apache.org/jaxws
    http://cxf.apache.org/schemas/jaxws.xsd">

  <jaxws:client name="{http://teiid.org}port_x"
    createdFromAPI="true">
    <jaxws:outInterceptors>
      <bean/>
      <ref bean="Timestamp_Request"/>
    </jaxws:outInterceptors>
  </jaxws:client>

  <bean
    id="Timestamp_Request">
    <constructor-arg>
      <map>
        <entry key="action" value="Timestamp"/>
      </map>
    </constructor-arg>
  </bean>

</beans>
```

Note that the client port configuration is matched to the data source instance by the QName `{http://teiid.org}port_x`, where the namespace will match your namespace setting or the default of `http://teiid.org`. The configuration may contain other port configurations with different local names.

WS-Security Kerberos is only supported when the WSDL property is defined in resource-adapter connection configuration and only when WSDL Based Procedures are used. WSDL file must contain WS-Policy section, then WS-Policy section is correctly interpreted and enforced on the endpoint. This is what the sample CXF configuration looks like:

```

<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:http="http://cxf.apache.org/transport/http/configuration"
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xmlns:cxf="http://cxf.apache.org/core"
  xmlns:p="http://cxf.apache.org/policy"
  xmlns:sec="http://cxf.apache.org/configuration/security"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
http://cxf.apache.org/jaxws
http://cxf.apache.org/schemas/jaxws.xsd
http://cxf.apache.org/transport/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
http://cxf.apache.org/configuration/security
http://cxf.apache.org/schemas/configuration/security.xsd
http://cxf.apache.org/core http://cxf.apache.org/schemas/core.xsd
http://cxf.apache.org/policy http://cxf.apache.org/schemas/policy.xsd">
  <bean
class="org.springframework.beans.factory.config.PropertyPlaceholderConfigu
rer"/>
  <cxf:bus>
    <cxf:features>
      <p:policies/>
      <cxf:logging/>
    </cxf:features>
  </cxf:bus>

  <jaxws:client name="
{http://webservices.samples.jboss.org/}HelloWorldPort"
createdFromAPI="true">
    <jaxws:properties>
      <entry key="ws-security.kerberos.client">
        <bean
class="org.apache.cxf.ws.security.kerberos.KerberosClient">
          <constructor-arg ref="cxf"/>
          <property name="contextName" value="alice"/>
          <property name="serviceName"
value="bob@service.example.com"/>
        </bean>
      </entry>
    </jaxws:properties>
  </jaxws:client>
</beans>

```

Nexst, configure the security-domain in the `standalone-teiid.xml` file under 'security' subsystem like this:

```

<security-domain name="alice" cache-type="default">
  <authentication>
    <login-module code="Kerberos" flag="required">
      <module-option name="storeKey" value="true"/>
      <module-option name="useKeyTab" value="true"/>
      <module-option name="keyTab"
value="/home/alice/alice.keytab"/>
    </login-module>
  </authentication>
</security-domain>

```

```

        <module-option name="principal" value="alice@EXAMPLE.COM"/>
        <module-option name="doNotPrompt" value="true"/>
        <module-option name="debug" value="true"/>
        <module-option name="refreshKrb5Config" value="true"/>
    </login-module>
</authentication>
</security-domain>

```

The CXF config property may also be used to control the logging of requests and responses for specific or all ports. Logging, when enabled, will be performed at an INFO level to the `org.apache.cxf.interceptor.context:batch /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS:add(jndi-name=java:/wsDS, class-name=org.teiid.resource.adapter.ws.WSManagedConnectionFactory, enabled=true, use-java-context=true) /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigFile:add(value=${jboss.server.home.dir}/standalone/configuration/xxx-jbossws-cxf.xml) /subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigName:add(value=port_x) /subsystem=resource-adapters/resource-adapter=webservice:activate runbatch`

Here is the corresponding `xxx-jbossws-cxf.xml` file:

```

<beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:jaxws="http://cxf.apache.org/jaxws"
    xsi:schemaLocation="http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans.xsd
        http://cxf.apache.org/jaxws
        http://cxf.apache.org/schemas/jaxws.xsd">

    <jaxws:client name="{http://teiid.org}port_y"
        createdFromAPI="true">
        <jaxws:features>
            <bean class="org.apache.cxf.feature.LoggingFeature"/>
        </jaxws:features>
    </jaxws:client>

</beans>

```

The CXF config property may also be used to control low level aspects of the HTTP transport:

```

<beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:http-
conf="http://cxf.apache.org/transports/http/configuration"

    xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
        http://cxf.apache.org/schemas/configuration/http-conf.xsd
        http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans.xsd">

```

```

<http-conf:conduit name="{http://teiid.org}port_z.http-conduit">
  <!-- WARNING ! disableCNcheck=true should NOT be used in production -
->
  <http-conf:tlsClientParameters disableCNcheck="true" />

</http-conf:conduit>
</beans>

```

To use HTTPS, configure the CXF file like this (You can also configure HTTPBasic, kerberos and so forth):

```

<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:sec="http://cxf.apache.org/configuration/security"
  xmlns:http-conf="http://cxf.apache.org/transports/http/configuration"
  xmlns:jaxws="http://java.sun.com/xml/ns/jaxws"

  xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
http://cxf.apache.org/configuration/security
http://cxf.apache.org/schemas/configuration/security.xsd">

  <http-conf:conduit name="*.http-conduit">
    <http-conf:client ConnectionTimeout="120000"
ReceiveTimeout="240000"/>
    <http-conf:tlsClientParameters secureSocketProtocol="SSL">
      <sec:trustManagers>
        <sec:keyStore type="JKS" password="changeit"
file="/path/to/truststore.jks"/>
      </sec:trustManagers>
    </http-conf:tlsClientParameters>
  </http-conf:conduit>
</beans>

```

The kerberos support is based SPNEGO as described in <http://cxf.apache.org/docs/client-http-transport-including-ssl-support.html#ClientHTTPTransport%28includingSSLsupport%29-SpnegoAuthentication%28Kerberos%29>. There two types of kerberos support, this first of which is Negotiation. With this configuration, the REST service is configured with a Kerberos JAAS domain, to negotiate a token, which is then used to access the web service. For this first create a security domain in the `standalone.xml` file:

```

<security-domain name="MY_REALM" cache-type="default">
  <authentication>
    <login-module code="Kerberos" flag="required">
      <module-option name="storeKey" value="true"/>

      <module-option name="useKeyTab" value="true"/>
      <module-option name="keyTab"
value="/home/username/service.keytab"/>
      <module-option name="principal"
value="host/testserver@MY_REALM"/>

      <module-option name="doNotPrompt" value="true"/>

```

```

        <module-option name="debug" value="false"/>
    </login-module>
</authentication>
</security-domain>

```

Set the `jboss-cxf-xxx.xml` file like this:

```

<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:sec="http://cxf.apache.org/configuration/security"
       xmlns:http-
conf="http://cxf.apache.org/transports/http/configuration"

xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
http://cxf.apache.org/configuration/security
http://cxf.apache.org/schemas/configuration/security.xsd">

    <http-conf:conduit name="*.http-conduit">
        <http-conf:authorization>
            <sec:AuthorizationType>Negotiate</sec:AuthorizationType>
            <sec:Authorization>MY_REALM</sec:Authorization>
        </http-conf:authorization>
    </http-conf:conduit>
</beans>


```

The resource adapter creation needs to define these properties:

```

<config-property name="ConfigFile">path/to/jboss-cxf-xxxx.xml</config-
property>
<config-property name="ConfigName">test</config-property>

```



IMPORTANT

Even though above configuration configures the value of "ConfigName", the cxf framework currently in the case of JAX-RS client does not give option to use it. For that reason use "*.http-conduit" which will apply to all the HTTP communications under this resource adapter.

The second method is Delegation. If the user is already logged into Teiid using Kerberos using JDBC/ODBC or used SPNEGO in web-tier and used pass-through authentication into Teiid, then there is no need to negotiate a new token for the Kerberos. The system can delegate the existing token. To configure for delegation, set up security domain defined exactly as defined in "negotiation", and `jboss-cxf-xxx.xml` file, however remove the following line from `jboss-cxf-xxx.xml` file, as it is not going to negotiate new token.

```

<sec:Authorization>MY_REALM</sec:Authorization>

```

Add the following properties in web service resource adapter creation. One determines that "kerberos" security being used, the second defines a security domain to be used at the data source, in this case we want to use a security domain that passes through a logged-in user:

-


```
<config-property name="SecurityType">Kerberos</config-property>
<security>
  <security-domain>passthrough-security</security-domain>
</security>
```

To configure the "passthrough-security" security domain you must add the "security" subsystem settings:

```
<security-domain name="passthrough-security" cache-type="default">
  <authentication>
    <login-module
code="org.teiid.jboss.PassthroughIdentityLoginModule" flag="required"
module="org.jboss.teiid">
      <module-option name="username" value="guest"/>
      <module-option name="password" value="guest"/>
    </login-module>
  </authentication>
</security-domain>
```



NOTE

The username and password are optional. If there is no authenticated subject available in the context, these can help create a simple static user, but that user will not work with kerberos authentication as the subject will not have the kerberos token that is required.

4.10. MANAGING DEPLOYED VIRTUAL DATABASES USING MANAGEMENT CONSOLE

4.10.1. Managing Deployed Virtual Databases Using Management Console

Administrators can manage deployed virtual databases from within the **Virtual Databases** panel.

4.10.2. Opening the Virtual Databases Panel

1. Log in to the JBoss Management Console.
2. Select the **Runtime** tab.
3. From the navigation tree, select **Status** → **Subsystems** → **Virtual Databases**.

CHAPTER 5. VERSIONING

5.1. VIRTUAL DATABASE VERSIONING

Virtual database (VDB) versioning allows you to deploy multiple versions of a VDB simultaneously, and to specify which version will be used in certain scenarios.

When an application connects to JBoss Data Virtualization, you can set the desired version of a VDB using a connection property. (Refer to the *Red Hat JBoss Data Virtualization User Guide* for more information.)

If a specific version is set, then you can only connect to that VDB. If no version is set, the deployed VDBs are searched until the appropriate version is found. This feature helps support more dynamic migration scenarios.

5.2. SET THE VDB VERSION

You can set the version in one of two ways: through the `vdb.xml` file, (which is useful for dynamic VDBs), or by specifying a naming convention in the deployment file (such as `VDBNAME.VERSION.vdb`). The deployer is responsible for choosing an appropriate version number. If there is already a VDB name and version combination that matches the current deployment, then connections to the previous VDB will be terminated and its cache entries will be flushed. Any new connections will then be made to the new VDB.

5.3. VIRTUAL DATABASE CONNECTION TYPE

Once your VDB is deployed, you can configure a property called connection type. Use this to determine what connections can be made to the VDB. The connection type can be one of:

- **NONE**: disallow new connections.
- **BY_VERSION**: (the default setting) allow connections only if the version is specified or if this is the earliest **BY_VERSION** VDB and there are no VDBs marked as **ANY**.
- **ANY**: allow connections with or without a version specified.

If you only want to migrate a few of your applications to the new version of the VDB, then set it to **BY_VERSION**. This ensures that only applications that know of the new version may use it.

If only a select few applications are to remain on the current VDB version, then you will need to update their connection settings to reference the current VDB by its version. The newly deployed VDB will then have its connection type set to **ANY**, which allows all new connections to be made against the newer version.

If you need to undertake a rollback in this scenario, then the newly-deployed VDB will, accordingly, have its connection type set to **NONE** or **BY_VERSION**.

5.4. SET THE VDB CONNECTION TYPE VIA ADMIN API

You can change a VDB connection type using the `changeVDBConnectionType` method provided by the Admin interface within the Admin API package (`org.teiid.adminapi`).

Javadocs for Red Hat JBoss Data Virtualization can be found on the [Red Hat Customer Portal](#).

CHAPTER 6. RESOURCE ADAPTERS

6.1. RESOURCE ADAPTERS IN JBOSS DATA VIRTUALIZATION

With the exception of JDBC data sources, JBoss Data Virtualization provides a JCA adapter for each supported data source. These are the resource adapter identifiers, as specified in the server configuration file:

- File Adapter - `file`
- Google Spreadsheet Adapter - `google`
- Red Hat JBoss Data Grid (6.1 & 6.2) Adapter - `infinispan`
- LDAP Adapter - `ldap`
- Salesforce Adapter - `salesforce`
- Web Services Adapter - `webservice`
- Mongo DB Adapter (technical preview) - `mongodb`




NOTE

A resource adapter for the JDBC translator is provided with JBoss EAP by default.


6.2. CONFIGURING RESOURCE ADAPTERS

You can use Management Console, AdminShell, or Management CLI (or, in standalone mode, directly edit the server configuration file) to configure resource adapters for data sources required by a VDB.



NOTE

Note that in domain mode, you must use Management CLI, Management Console or AdminShell to configure data sources. Refer to the JBoss Enterprise Application Platform *Administration and Configuration Guide* for an example of how to configure resource adapters using the Management Console.



IMPORTANT

If you configure a resource adapter using either a CLI script or the AdminShell, the name of the resource adapter must not be the same as that of an existing one. (Teiid Designer handles this by creating a unique name for the resource adapter.)

The reasons for this are as follows:

- Any time a resource adapter with a duplicate name is added, the server has to be restarted. This conforms with the way Red Hat JBoss EAP's JCA system was designed, in that multiple instances of resource adapters with the same name are only recognised when the server is launched.
- Anytime you delete a resource adapter by name, all instances with the same name will be deleted.

6.3. EXAMPLE CONFIGURATION

Example configuration, as it would appear in the server configuration file, can be found in the `EAP_HOME/docs/teiid/datasources/` directory.



NOTE

This configuration will need to be adjusted with file paths and properties appropriate to your installation. The JNDI name must be the same JNDI name used in the VDB.

6.4. RESOURCE ADAPTER PROPERTIES

For a full listing of configuration properties, you can run the following command from within the Management CLI:

```
/subsystem=teiid:read-rar-description(rar-name=ADAPTER_ID)
```

`ADAPTER_ID` refers to the identifier of the resource adapter as specified in the server configuration file.

6.5. CONFIGURING RESOURCE ADAPTERS USING CLI SCRIPTS

This is the preferred method for configuring resource adapters.

JBoss Data Virtualization provides an example Management CLI script for configuring each of the provided resource adapters. A script for a particular resource adapter can be found in the `EAP_HOME/docs/teiid/datasources/DATASOURCE` directory.

When configuring resource adapters using the Management CLI, you must provide a properties file. Sample properties files are provided in the same directory as the sample scripts.

You can run these scripts (once JBoss Data Virtualization is running) by executing the following command:

```
./EAP_HOME/bin/jboss-cli.sh --connect --
file=EAP_HOME/docs/teiid/datasources/DATASOURCE/SCRIPT.cli --
properties=EAP_HOME/docs/teiid/datasources/DATASOURCE/SCRIPT.properties
```

For more information about using the Management CLI, see the JBoss Enterprise Application Platform *Administration and Configuration Guide*



NOTE

These scripts will need to be adjusted with file paths and properties appropriate to your installation. The JNDI name must be the same JNDI name used in the VDB.

6.6. FILE ADAPTER PROPERTIES

The following table describes the configuration properties that can be configured for the File resource adapter:

Config property	Example	Description
ParentDirectory		Directory where the data files are stored.
FileMapping	file1.txt=fileX.txt,file2.txt=fileY.txt	Set FileMapping to redirect specific relative paths (case sensitive) to alternative locations. The string value specifies a map in the format key=value(,key=value)*. Optional.
AllowParentPaths	true	Set AllowParentPaths to false to disallow '..' in paths. This prevents requesting files that are not contained in the parent directory. Optional.

6.7. GOOGLE SPREADSHEET RESOURCE ADAPTER PROPERTIES

6.7.1. Google Spreadsheet Resource Adapter Properties

The following table describes the configuration properties that can be configured for the Google Spreadsheet resource adapter:

Config property	Description
AuthMethod	This is the authentication method used to access Google. This property can be set to either ClientLogin or OAuth2 . If the setting is ClientLogin it is necessary to provide two additional properties: Username and Password . If the setting is OAuth2 it is necessary to provide a RefreshToken .
RefreshToken	This is required only if AuthMethod=OAuth2
Username	Username for the Google account. Required only if AuthMethod=ClientLogin
Password	Password for the Google account. Required only if AuthMethod=ClientLogin
SpreadsheetName	The name of the spreadsheet to which this resource adapter is connecting. Required.
BatchSize	The maximum number of rows that can be fetched at a time. Default is 4096.

6.7.2. Obtaining an OAuth Refresh Token

When using the Google Spreadsheet resource adapter with OAuth authentication, you will need to obtain an OAuth refresh token.

Procedure 6.1. Obtaining an OAuth Refresh Token

1. Get an authorization code

Click on the following link: [Get Authorization Code](#)

Click on **Allow** access to allow the **Teiid Google Connector** to access the Google account in which the spreadsheet resides.

2. Obtain the refresh token

Copy the authorization code from the previous step into the `code` field of the following POST request and run it from the command line:

```
curl --data-urlencode code=AUTH_CODE \
--data-urlencode client_id=217138521084.apps.googleusercontent.com \
--data-urlencode client_secret=gXQ6-l0kEjE1lVcz7giB4Poy \
--data-urlencode redirect_uri=urn:ietf:wg:oauth:2.0:oob \
--data-urlencode grant_type=authorization_code
https://accounts.google.com/o/oauth2/token
```

The refresh token will be in the response.

6.8. JBOSS DATA GRID RESOURCE ADAPTER PROPERTIES

The JBoss Data Grid (JDG) resource adapter can be configured to support the following caching modes:

Cache Type	Obtain Cache By
Remote Cache	using JNDI
Remote Cache	1 or more host:port combinations specified
Remote Cache	referring to HotRod client properties file specified

The following table describes the configuration properties that can be configured for the resource adapter:

Property Name	Required	Property Template	Description
CacheTypeMap	Y	cacheName:className[; pkFieldName] [,cacheName:className ;pkFieldName..]	This property maps the root Java class name to the cache and identifies the primary key.

Property Name	Required	Property Template	Description
module	N		This property specifies the JBoss EAP module containing the cache classes defined in CacheTypeMap.
CacheJndiName	N		This is the JNDI name used to find the CacheContainer.
RemoteServerList	N	host:port[;host:portâ;.]	This property specifies the host (and ports) that will be clustered together to access the caches defined in CacheTypeMap.
ConfigurationFileNameForLocalCache	N		This is the XML configuration file for configuring a local cache.
HotRodClientPropertiesFile	N		This is the HotRod properties file for configuring a connection to a remote cache.

6.9. LDAP ADAPTER PROPERTIES

The following table describes the configuration properties that can be configured for the LDAP resource adapter:

Config property	Property Template	Description
LdapUrl		LDAP Directory URL. This property is mandatory.
LdapAdminUserDN	cn=???,ou=???,dc=???	LDAP administration user DN. This property is mandatory.
LdapAdminUserPassword		LDAP administration password. This property is mandatory.
LdapTxnTimeoutInMillis		LDAP transaction timeout in milliseconds. -1 = no timeout. This property is optional.

6.10. SALESFORCE ADAPTER PROPERTIES

The following table describes the configuration properties that can be configured for the Salesforce resource adapter:

Config property	Description
URL	The URL to connect to.
username	The username.
password	The password.
requestTimeout	This is an optional property for setting timeouts, which can also be done through the CXF config.
connectTimeout	This is an optional property for setting timeouts, which can also be done through the CXF config.
configFile	Use this property to supply specific configuration for the Salesforce service. This configuration must contain config for "SforceService" service with namespace "urn:partner.soap.sforce.com".

6.11. WEB SERVICES ADAPTER PROPERTIES

The following table describes the configuration properties that can be configured for the Web Services resource adapter:

Config property	Example	Description
EndPoint		End point URL for the web service
SecurityType	HTTPBasic	Use for http basic security.
AuthUserName		Use for http basic security.
AuthPassword		Use for http basic security.
RequestTimeout		This is an optional property for setting timeouts, which can also be done through the CXF config.
ConnectTimeout		This is an optional property for setting timeouts, which can also be done through the CXF config.
ConfigFile		Use these properties to supply specific CXF configuration for this service. This file must contain a configuration for the name defined on the "EndPointName" property.

Config property	Example	Description
EndPointName	WebSVC	Use with ConfigFile. These properties to supply specific CXF configuration for this service.

6.12. MONGODB ADAPTER PROPERTIES

The following table describes the configuration properties that can be configured for the MongoDB resource adapter:

Config property	Example	Description
RemoteServerList	localhost:27017	MongoDB server list in this form host:port[;host:port...]*
Database		Database name.
Username		Use this property together with Password to supply credentials.
Password		Use this property together with Username to supply credentials.

CHAPTER 7. CXF CONFIGURATION

7.1. CXF CONFIGURATION

If your configuration uses CXF, you need to specify the data source configuration file and port configuration.

The `ConfigFile` property, set in the server configuration file for the resource-adapter subsystem, specifies the Spring XML configuration file. (The system's default configuration is used if the configuration file is not specified.)

The `ConfigName` property, also set in the server configuration file for the resource-adapter subsystem, specifies the port configuration. Set this to the local part of the desired port QName.

The configuration file property is specified via `/subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigFile:add(value=${jboss.server.home.dir}/standalone/configuration/xxx-jbossws-cxf.xml)`

7.2. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE

Prerequisites

- You must have configured a web service data source.

Procedure 7.1. Configure CXF for a Web Service Data Source

1. Specify the Web Service CXF `ConfigFile`

Run the following command from within the Management CLI, specifying the CXF configuration file for the data source:

```
/subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigFile:add(value=CONFIG-FILE.xml)
```

2. Specify the Web Service CXF `ConfigName`

Specify the port configuration by running the following command from within the Management CLI, using the port QName (local part only) for the value:

```
/subsystem=resource-adapters/resource-adapter=webservice/connection-definitions=wsDS/config-properties=ConfigName:add(value=CONFIG-NAME)
```

3. Create/Edit the CXF Configuration File

Open/create the `EAP_HOME/MODE/configuration/CONFIG-FILE.xml` configuration file.

```
<http-conf:conduit name="{NAMESPACE}CONFIG-NAME.http-conduit">
  ...
</http-conf:conduit>
```

**NOTE**


CONFIG-NAME is the same as that specified above, with a default value of `teiid`. *NAMESPACE* is the namespace URI for the QName in your config file, which should match your WSDL/namespace setting on the data source or use the default of "http://teiid.org". The namespace may be set via the namespace datasource property. Typically that will only need done when also supplying the WSDL setting.

The following is an example of a CXF file configuring timeouts:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:http-
conf="http://cxf.apache.org/transport/http/configuration"

xsi:schemaLocation="http://cxf.apache.org/transport/http/figurati
ion
  http://cxf.apache.org/schemas/configuration/http-conf.xsd
  http://www.springframework.org/schema/beans
  http://www.springframework.org/schema/beans/spring-beans.xsd">

  <http-conf:conduit name="{NAMESPACE}CONFIG-NAME.http-conduit">
    <http-conf:client ConnectionTimeout="120000"
ReceiveTimeout="240000"/>
  </http-conf:conduit>
</beans>
```

**NOTE**

For web service data sources, CXF configuration is only applicable to non-binary calls.

7.3. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: WS-SECURITY

Prerequisites

- The web service data source must be configured and the `ConfigFile` and `ConfigName` properties must be configured for CXF.

Procedure 7.2. Configure CXF for a Web Service Data Source: WS-Security

1. Specify the CXF SecurityType

Run the following command from within the Management CLI, using `WSecurity` as the value for `SecurityType`:

```
/subsystem=resource-adapters/resource-adapter=webservice/connection-
definitions=wsDS/config-
properties=SecurityType:add(value=WSecurity)
```

2. Modify the CXF Configuration File

Open the CXF configuration file for the web service data source and add your desired properties.

The following is an example of a web service data source CXF configuration file adding a timestamp to the SOAP header:


```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:jaxws="http://cxf.apache.org/jaxws"

       xsi:schemaLocation="http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-
                           beans.xsd
                           http://cxf.apache.org/jaxws
                           http://cxf.apache.org/schemas/jaxws.xsd">

       <jaxws:client name="{http://teiid.org}.teiid"
                   createdFromAPI="true">
         <jaxws:outInterceptors>
           <ref bean="Timestamp_Request"/>
         </jaxws:outInterceptors>
       </jaxws:client>

       <bean
         id="Timestamp_Request">
         <constructor-arg>
           <map>
             <entry key="action" value="Timestamp"/>
           </map>
         </constructor-arg>
       </bean>

</beans>
```



NOTE

- A WSDL is not expected to describe the service being used.
- The Spring XML configuration file must contain the relevant policy configuration.
- The client port configuration is matched to the data source instance by the **CONFIG-NAME**. The configuration may contain other port configurations with different local names.

References

- For more information about WS-Security and CXF configuration options refer to <http://cxf.apache.org/docs/ws-security.html>.

7.4. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: LOGGING

CXF configuration can control the logging of requests and responses for specific or all ports. Logging, when enabled, is performed at an **INFO** level to the `org.apache.cxf.interceptor` context.

Prerequisites

- The web service data source must be configured and the `ConfigFile` and `ConfigName` properties must be configured for CXF.

Procedure 7.3. Configure CXF for a Web Service Data Source: Logging

- **Modify the CXF Configuration File**

Open the CXF configuration file for the web service data source and add your desired logging properties.

The following is an example of a CXF configuration file for a web service data source that enables logging:

```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:jaxws="http://cxf.apache.org/jaxws"

       xsi:schemaLocation="http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-
                           beans.xsd
                           http://cxf.apache.org/jaxws
                           http://cxf.apache.org/schemas/jaxws.xsd">

       <jaxws:client name="{http://teiid.org}teiid"
                   createdFromAPI="true">
         <jaxws:features>
           <bean class="org.apache.cxf.feature.LoggingFeature"/>
         </jaxws:features>
       </jaxws:client>

</beans>
```

References

- For more information about CXF logging configuration options see <http://cxf.apache.org/docs/debugging-and-logging.html>.

7.5. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: TRANSPORT SETTINGS

CXF configuration can also control low level aspects of the HTTP transport.

Prerequisites

- The web service data source must be configured and the `ConfigFile` and `ConfigName` properties must be configured for CXF.

Procedure 7.4. Configure CXF for a Web Service Data Source: Transport Settings

- Open the CXF configuration file for the web service data source and add your desired transport properties.

The following is an example of a CXF configuration file for a web service data source that disables hostname verification:

```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:jaxws="http://cxf.apache.org/jaxws"

       xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
                           http://cxf.apache.org/schemas/configuration/http-conf.xsd
                           http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-
                           beans.xsd">

       <http-conf:conduit name="{http://teiid.org}teiid.http-conduit">
         <http-conf:tlsClientParameters disableCNcheck="tru" />
       </http-conf:conduit>

</beans>
```



WARNING

`disableCNcheck=true` must NOT be used in production.

7.6. CONFIGURE CXF FOR A WEB SERVICE DATA SOURCE: SSL SUPPORT (HTTPS)

For using HTTPS, you can configure the CXF configuration file as below:

```
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:sec="http://cxf.apache.org/configuration/security"
       xmlns:http-conf="http://cxf.apache.org/transports/http/configuration"
       xmlns:jaxws="http://java.sun.com/xml/ns/jaxws"

       xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
                           http://cxf.apache.org/schemas/configuration/http-conf.xsd
                           http://www.springframework.org/schema/beans
                           http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
                           http://cxf.apache.org/configuration/security
                           http://cxf.apache.org/schemas/configuration/security.xsd">

       <http-conf:conduit name="*.http-conduit">
         <http-conf:client ConnectionTimeout="120000"
           ReceiveTimeout="240000"/>
         <http-conf:tlsClientParameters secureSocketProtocol="SSL">
           <sec:trustManagers>
             <sec:keyStore type="JKS" password="changeit"
               file="/path/to/truststore.jks"/>
           </sec:trustManagers>
         </http-conf:tlsClientParameters>
       </http-conf:conduit>

</beans>
```

```

        </sec:trustManagers>
    </http-conf:tlsClientParameters>
</http-conf:conduit>
</beans>

```

For more information about http-conduit based configuration see <http://cxf.apache.org/docs/client-http-transport-including-ssl-support.html>. You can also configure for services such as HTTPBasic and Kerberos.

7.7. CONFIGURE CXF FOR A SALESFORCE DATA SOURCE

Prerequisites

- You must have configured a Salesforce data source.

Procedure 7.5. Configure CXF for a Salesforce Data Source

1. Run the following command from within the Management CLI, specifying the CXF configuration file for the data source:

```

/subsystem=resource-adapters/resource-adapter=salesforce/connection-
definitions=sfDS/config-properties=ConfigFile:add(value=CONFIG-
FILE.xml)

```

2. Specify the port configuration by running the following command from within the Management CLI, using the port QName (local part only) for the value, in this case **Soap**:

```

/subsystem=resource-adapters/resource-adapter=salesforce/connection-
definitions=sfDS/config-properties=ConfigName:add(value=Soap)

```

3. Open/create the **EAP_HOME/MODE/configuration/CONFIG-FILE.xml** configuration file. Set the namespace URI for the QName to **{urn:partner.soap.sforce.com}**, using **Soap** as the value for **ConfigName**:

```

<http-conf:conduit name="{urn:partner.soap.sforce.com}Soap.http-
conduit">
...
</http-conf:conduit>

```

The following is an example of a CXF file that configures timeout values:

```


<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:http-
conf="http://cxf.apache.org/transports/http/configuration"

  xsi:schemaLocation="http://cxf.apache.org/transports/http/configurat
ion
  http://cxf.apache.org/schemas/configuration/http-conf.xsd
  http://www.springframework.org/schema/beans
  http://www.springframework.org/schema/beans/spring-beans.xsd">

  <http-conf:conduit name="{urn:partner.soap.sforce.com}Soap.http-

```

```
conduit">  
  <http-conf:client ConnectionTimeout="120000"  
ReceiveTimeout="240000"/>  
</http-conf:conduit>  
</beans>
```



NOTE

CXF configuration for Salesforce data sources is only used for http bus configuration. It is not used for WS-Security. Salesforce has its own security authentication.

CHAPTER 8. JDBC DATA SOURCES

8.1. CONFIGURING JDBC DATA SOURCES

For JBoss Data Virtualization to connect to JDBC data sources, you must configure the data source in the JBoss EAP instance with the following two steps:

1. Install (or deploy) the JDBC JAR driver file.
2. Create (or configure) it as a data source in JBoss EAP.



NOTE

See the *Red Hat JBoss Enterprise Application Platform Administration and Configuration Guide* for more comprehensive information about configuring data sources.

8.2. EXAMPLE CONFIGURATION

Example configuration, as it would appear in the server configuration file, can be found in the `EAP_HOME/docs/teiid/datasources/` directory.



NOTE

This configuration will need to be adjusted with file paths and properties appropriate to your installation. The JNDI name must be the same JNDI name used in the VDB.

8.3. INSTALL A JDBC DRIVER WITH THE MANAGEMENT CLI

Procedure 8.1. Install a JDBC Driver with the Management CLI

1. Start the Management CLI: `./EAP_HOME/bin/jboss-cli.sh`
2. Enter the `connect` command.
3. Enter the `deploy PATH/FILE.jar` command: `deploy ojdbc6.jar`
4. Enter the `quit` command.

8.4. INSTALL A JDBC DRIVER WITH THE MANAGEMENT CONSOLE

Summary

Before your application can connect to a JDBC datasource, your datasource vendor's JDBC drivers need to be installed in a location where JBoss EAP 6 can use them. JBoss EAP 6 allows you to deploy these drivers like any other deployment. This means that you can deploy them across multiple servers in a server group, if you use a managed domain.

Prerequisites

Before performing this task, you need to meet the following prerequisites:

- Download the JDBC driver from your database vendor.

**NOTE**

Any JDBC 4-compliant driver is automatically recognized and installed in the system by name and version. A JDBC JAR is identified using the Java service provider mechanism. Such JARs contain the META-INF/services/java.sql.Driver text, which contains the name of the Driver classes in that JAR.

Procedure 8.2. Modify the JDBC Driver JAR

If the JDBC driver JAR is not JDBC 4-compliant, it can be made deployable using the following method.

1. Change to, or create, an empty temporary directory.
2. Create a META-INF subdirectory.
3. Create a META-INF/services subdirectory.
4. Create a META-INF/services/java.sql.Driver file, which contains one line indicating the fully-qualified class name of the JDBC driver.
5. Use the JAR command-line tool to update the JAR like this:

```
jar \-uf jdbc-driver.jar META-INF/services/java.sql.Driver
```

6. [Section 2.2.2, “Log in to the Management Console”](#)
7. If you use a managed domain, deploy the JAR file to a server group. Otherwise, deploy it to your server. See the Deploy with the Management Console section of the Administration and Configuration Guide for JBoss EAP 6 for more information.

Result:

The JDBC driver is deployed, and is available for your applications to use.

8.5. CREATE A NON-XA DATASOURCE WITH THE MANAGEMENT INTERFACES

Summary

This topic covers the steps required to create a non-XA datasource, using either the Management Console or the Management CLI.

Prerequisites

- The JBoss EAP 6 server must be running.

**NOTE**

Prior to version 10.2 of the Oracle datasource, the `<no-tx-separate-pools/>` parameter was required, as mixing non-transactional and transactional connections would result in an error. This parameter may no longer be required for certain applications.

**NOTE**

To prevent issues such as duplication of driver listing, selected driver not available in a profile, or driver not displayed if a server for the profile is not running, in JBoss EAP 6.4 onwards, only JDBC drivers that are installed as modules and correctly referenced from profiles are detectable while creating a datasource using the Management Console in domain mode.

Procedure 8.3. Create a Datasource using either the Management CLI or the Management Console

- ○ **Management CLI**

- a. Launch the CLI tool and connect to your server.
- b. Run the following Management CLI command to create a non-XA datasource, configuring the variables as appropriate:

**NOTE**

The value for *DRIVER_NAME* depends on the number of classes listed in the `/META-INF/services/java.sql.Driver` file located in the JDBC driver JAR. If there is only one class, the value is the name of the JAR. If there are multiple classes, the value is the name of the JAR + `driverClassName + "_" + majorVersion + "_" + minorVersion`. Failure to do so will result in the following error being logged:

```
JBAS014775:      New missing/unsatisfied dependencies
```

For example, the *DRIVER_NAME* value required for the MySQL 5.1.31 driver, is `mysql-connector-java-5.1.31-bin.jarcom.mysql.jdbc.Driver_5_1`.

```
data-source add --name=DATASOURCE_NAME --jndi-name=JNDI_NAME -
-driver-name=DRIVER_NAME --connection-url=CONNECTION_URL
```

- c. Enable the datasource:

```
data-source enable --name=DATASOURCE_NAME
```

- **Management Console**

- a. Login to the Management Console.
- b. **Navigate to the Datasources panel in the Management Console**
 - i. Select the **Configuration** tab from the top of the console.
 - ii. For Domain mode only, select a profile from the drop-down box in the top left.
 - iii. Expand the **Subsystems** menu on the left of the console, then expand the **Connector** menu.
 - iv. Select **Datasources** from the menu on the left of the console.

c. Create a new datasource

- i. Click **Add** at the top of the **Datasources** panel.
- ii. Enter the new datasource attributes in the **Create Datasource** wizard and proceed with the **Next** button.
- iii. Enter the JDBC driver details in the **Create Datasource** wizard and click **Next** to continue.
- iv. Enter the connection settings in the **Create Datasource** wizard.
- v. Click the **Test Connection** button to test the connection to the datasource and verify the settings are correct.
- vi. Click **Done** to finish

Result

The non-XA datasource has been added to the server. It is now visible in either the `standalone.xml` or `domain.xml` file, as well as the management interfaces.

8.6. CREATE AN XA DATASOURCE WITH THE MANAGEMENT INTERFACES

Summary

This topic covers the steps required to create an XA datasource, using either the Management Console or the Management CLI.



NOTE

Prior to version 10.2 of the Oracle datasource, the `<no-tx-separate-pools/>` parameter was required, as mixing non-transactional and transactional connections would result in an error. This parameter may no longer be required for certain applications.

Procedure 8.4. Create an XA Datasource, Using Either the Management CLI or the Management Console

- ○ **Management CLI**
 - a. [Section 2.3.1, “Launch the Management CLI”](#).
 - b. Run the following Management CLI command to create an XA datasource, configuring the variables as appropriate:

**NOTE**

The value for *DRIVER_NAME* depends on the number of classes listed in the `/META-INF/services/java.sql.Driver` file located in the JDBC driver JAR. If there is only one class, the value is the name of the JAR. If there are multiple classes, the value is the name of the JAR + `driverClassName + "_" + majorVersion + "_" + minorVersion`. Failure to do so will result in the following error being logged:

```
JBAS014775: New missing/unsatisfied dependencies
```

For example, the *DRIVER_NAME* value required for the MySQL 5.1.31 driver, is `mysql-connector-java-5.1.31-bin.jarcom.mysql.jdbc.Driver_5_1`.

```
xa-data-source add --name=XA_DATASOURCE_NAME --jndi-name=JNDI_NAME --driver-name=DRIVER_NAME --xa-datasource-class=XA_DATASOURCE_CLASS
```

c. Configure the XA datasource properties**i. Set the server name**

Run the following command to configure the server name for the host:

```
/subsystem=datasources/xa-data-source=XA_DATASOURCE_NAME/xa-datasource-properties=ServerName:add(value=HOSTNAME)
```

ii. Set the database name

Run the following command to configure the database name:

```
/subsystem=datasources/xa-data-source=XA_DATASOURCE_NAME/xa-datasource-properties=DatabaseName:add(value=DATABASE_NAME)
```

d. Enable the datasource:

```
xa-data-source enable --name=XA_DATASOURCE_NAME
```

o Management Console**a. [Section 2.2.2, “Log in to the Management Console”](#).****b. Navigate to the Datasources panel in the Management Console**

i. Select the **Configuration** tab from the top of the console.

ii. For Domain mode only, select a profile from the drop-down box at the top left.

iii. Expand the **Subsystems** menu on the left of the console, then expand the **Connector** menu.

iv. Select **Datasources**.

- c. Select the **XA Datasource** tab.
- d. **Create a new XA datasource**
 - i. Click **Add**.
 - ii. Enter the new XA datasource attributes in the **Create XA Datasource** wizard and click **Next**.
 - iii. Enter the JDBC driver details in the **Create XA Datasource** wizard and click **Next**.
 - iv. Enter the XA properties and click **Next**.
 - v. Enter the connection settings in the **Create XA Datasource** wizard.
 - vi. Click the **Test Connection** button to test the connection to the XA datasource and verify the settings are correct.
 - vii. Click **Done** to finish

Result

The XA datasource has been added to the server. It is now visible in either the `standalone.xml` or `domain.xml` file, as well as the management interfaces.

8.7. DATASOURCE PARAMETERS

Table 8.1. Datasource parameters common to non-XA and XA datasources

Parameter	Description
<code>jndi-name</code>	The unique JNDI name for the datasource.
<code>pool-name</code>	The name of the management pool for the datasource.
<code>enabled</code>	Whether or not the datasource is enabled.
<code>use-java-context</code>	Whether to bind the datasource to global JNDI.
<code>spy</code>	Enable spy functionality on the JDBC layer. This logs all JDBC traffic to the datasource. Note that the logging category <code>jboss.jdbc.spy</code> must also be set to the log level DEBUG in the logging subsystem.
<code>use-ccm</code>	Enable the cached connection manager.
<code>new-connection-sql</code>	A SQL statement which executes when the connection is added to the connection pool.

Parameter	Description
transaction-isolation	One of the following: <ul style="list-style-type: none"> TRANSACTION_READ_UNCOMMITTED TRANSACTION_READ_COMMITTED TRANSACTION_REPEATABLE_READ TRANSACTION_SERIALIZABLE TRANSACTION_NONE
url-selector-strategy-class-name	A class that implements interface <code>org.jboss.jca.adapters.jdbc.URLSelectorStrategy</code> .
security	Contains child elements which are security settings. See Table 8.6, “Security parameters” .
validation	Contains child elements which are validation settings. See Table 8.7, “Validation parameters” .
timeout	Contains child elements which are timeout settings. See Table 8.8, “Timeout parameters” .
statement	Contains child elements which are statement settings. See Table 8.9, “Statement parameters” .

Table 8.2. Non-XA datasource parameters

Parameter	Description
jta	Enable JTA integration for non-XA datasources. Does not apply to XA datasources.
connection-url	The JDBC driver connection URL.
driver-class	The fully-qualified name of the JDBC driver class.
connection-property	Arbitrary connection properties passed to the method <code>Driver.connect(url, props)</code> . Each connection-property specifies a string name/value pair. The property name comes from the name, and the value comes from the element content.
pool	Contains child elements which are pooling settings. See Table 8.4, “Pool parameters common to non-XA and XA datasources” .
url-delimiter	The delimiter for URLs in a connection-url for High Availability (HA) clustered databases.

Table 8.3. XA datasource parameters

Parameter	Description
xa-datasource-property	A property to assign to implementation class <code>XADatasource</code> . Specified by <code>name=value</code> . If a setter method exists, in the format <code>setName</code> , the property is set by calling a setter method in the format of <code>setName(value)</code> .
xa-datasource-class	The fully-qualified name of the implementation class <code>javax.sql.XADatasource</code> .
driver	A unique reference to the class loader module which contains the JDBC driver. The accepted format is <code>driverName#majorVersion.minorVersion</code> .
xa-pool	Contains child elements which are pooling settings. See Table 8.4, “Pool parameters common to non-XA and XA datasources” and Table 8.5, “XA pool parameters”.
recovery	Contains child elements which are recovery settings. See Table 8.10, “Recovery parameters”.

Table 8.4. Pool parameters common to non-XA and XA datasources

Parameter	Description
min-pool-size	The minimum number of connections a pool holds.
max-pool-size	The maximum number of connections a pool can hold.
prefill	Whether to try to prefill the connection pool. The default is <code>false</code> .
use-strict-min	Whether the idle connection scan should strictly stop marking for closure of any further connections, once the <code>min-pool-size</code> has been reached. The default value is <code>false</code> .
flush-strategy	Whether the pool is flushed in the case of an error. Valid values are: <ul style="list-style-type: none"> • <code>FailingConnectionOnly</code> • <code>IdleConnections</code> • <code>EntirePool</code> The default is <code>FailingConnectionOnly</code> .

Parameter	Description
allow-multiple-users	Specifies if multiple users will access the datasource through the <code>getConnection(user, password)</code> method, and whether the internal pool type accounts for this behavior.

Table 8.5. XA pool parameters

Parameter	Description
is-same-rm-override	Whether the <code>javax.transaction.xa.XAResource.isSameRM(XAResource)</code> class returns <code>true</code> or <code>false</code> .
interleaving	Whether to enable interleaving for XA connection factories.
no-tx-separate-pools	Whether to create separate sub-pools for each context. This is required for Oracle datasources, which do not allow XA connections to be used both inside and outside of a JTA transaction. Using this option will cause your total pool size to be twice max-pool-size , because two actual pools will be created.
pad-xid	Whether to pad the Xid.
wrap-xa-resource	Whether to wrap the XAResource in an <code>org.jboss.tm.XAResourceWrapper</code> instance.

Table 8.6. Security parameters

Parameter	Description
user-name	The username to use to create a new connection.
password	The password to use to create a new connection.
security-domain	Contains the name of a JAAS security-manager which handles authentication. This name correlates to the <code>application-policy/name</code> attribute of the JAAS login configuration.
reauth-plugin	Defines a reauthentication plug-in to use to reauthenticate physical connections.

Table 8.7. Validation parameters

Parameter	Description
valid-connection-checker	An implementation of interface <code>org.jboss.jca.adapters.jdbc.ValidConnectionChecker</code> which provides a <code>SQLException.isValidConnection(Connection e)</code> method to validate a connection. An exception means the connection is destroyed. This overrides the parameter <code>check-valid-connection-sql</code> if it is present.
check-valid-connection-sql	An SQL statement to check validity of a pool connection. This may be called when a managed connection is taken from a pool for use.
validate-on-match	Indicates whether connection level validation is performed when a connection factory attempts to match a managed connection for a given set. Specifying "true" for <code>validate-on-match</code> is typically not done in conjunction with specifying "true" for <code>background-validation</code> . <code>Validate-on-match</code> is needed when a client must have a connection validated prior to use. This parameter is false by default.
background-validation	Specifies that connections are validated on a background thread. Background validation is a performance optimization when not used with <code>validate-on-match</code> . If <code>validate-on-match</code> is true, using <code>background-validation</code> could result in redundant checks. Background validation does leave open the opportunity for a bad connection to be given to the client for use (a connection goes bad between the time of the validation scan and prior to being handed to the client), so the client application must account for this possibility.
background-validation-millis	The amount of time, in milliseconds, that background validation runs.
use-fast-fail	If true, fail a connection allocation on the first attempt, if the connection is invalid. Defaults to <code>false</code> .
stale-connection-checker	An instance of <code>org.jboss.jca.adapters.jdbc.StaleConnectionChecker</code> which provides a Boolean <code>isStaleConnection(SQLException e)</code> method. If this method returns <code>true</code> , the exception is wrapped in an <code>org.jboss.jca.adapters.jdbc.StaleConnectionException</code> , which is a subclass of <code>SQLException</code> .

Parameter	Description
exception-sorter	An instance of <code>org.jboss.jca.adapters.jdbc.ExceptionSorter</code> which provides a Boolean <code>isExceptionFatal(SQLException e)</code> method. This method validates whether an exception is broadcast to all instances of <code>javax.resource.spi.ConnectionEventListener</code> as a <code>connectionErrorOccurred</code> message.

Table 8.8. Timeout parameters

Parameter	Description
use-try-lock	Uses <code>tryLock()</code> instead of <code>lock()</code> . This attempts to obtain the lock for the configured number of seconds, before timing out, rather than failing immediately if the lock is unavailable. Defaults to 60 seconds. As an example, to set a timeout of 5 minutes, set <code><use-try-lock>300</use-try-lock></code> .
blocking-timeout-millis	The maximum time, in milliseconds, to block while waiting for a connection. After this time is exceeded, an exception is thrown. This blocks only while waiting for a permit for a connection, and does not throw an exception if creating a new connection takes a long time. Defaults to 30000, which is 30 seconds.
idle-timeout-minutes	The maximum time, in minutes, before an idle connection is closed. If not specified, the default is 30 minutes. The actual maximum time depends upon the <code>idleRemover</code> scan time, which is half of the smallest <code>idle-timeout-minutes</code> of any pool.
set-tx-query-timeout	Whether to set the query timeout based on the time remaining until transaction timeout. Any configured query timeout is used if no transaction exists. Defaults to false .
query-timeout	Timeout for queries, in seconds. The default is no timeout.
allocation-retry	The number of times to retry allocating a connection before throwing an exception. The default is 0 , so an exception is thrown upon the first failure.

Parameter	Description
allocation-retry-wait-millis	How long, in milliseconds, to wait before retrying to allocate a connection. The default is 5000, which is 5 seconds.
xa-resource-timeout	If non-zero, this value is passed to method <code>XAResource.setTransactionTimeout</code> .

Table 8.9. Statement parameters

Parameter	Description
track-statements	<p>Whether to check for unclosed statements when a connection is returned to a pool and a statement is returned to the prepared statement cache. If false, statements are not tracked.</p> <p>Valid values</p> <ul style="list-style-type: none"> • true: statements and result sets are tracked, and a warning is issued if they are not closed. • false: neither statements or result sets are tracked. • nowarn: statements are tracked but no warning is issued. This is the default.
prepared-statement-cache-size	The number of prepared statements per connection, in a Least Recently Used (LRU) cache.
share-prepared-statements	Whether JBoss EAP should cache, instead of close or terminate, the underlying physical statement when the wrapper supplied to the application is closed by application code. The default is false .

Table 8.10. Recovery parameters

Parameter	Description
recover-credential	A username/password pair or security domain to use for recovery.
recover-plugin	An implementation of the <code>org.jboss.jca.core.spi.recoveryRecoveryPlugin</code> class, to be used for recovery.

8.8. JDBC DRIVER DOWNLOAD LOCATIONS

The following table gives the standard download locations for JDBC drivers of common databases used with JBoss EAP 6.



NOTE

These links point to third-party websites which are not controlled or actively monitored by Red Hat. For the most up-to-date drivers for your database, check your database vendor's documentation and website.

Table 8.11. JDBC driver download locations

Vendor	Download Location
MySQL	http://www.mysql.com/products/connector/
PostgreSQL	http://jdbc.postgresql.org/
Oracle	http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html
IBM	http://www-306.ibm.com/software/data/db2/java/
Sybase	http://www.sybase.com/products/allproductsa-z/softwaredeveloperkit/jconnect
Microsoft	http://msdn.microsoft.com/data/jdbc/

CHAPTER 9. LOGGING

9.1. OVERVIEW OF LOGGING

JBoss EAP 6 provides highly configurable logging facilities for both its own internal use and for use by deployed applications. The logging subsystem is based on JBoss LogManager and it supports several third party application logging frameworks in addition to JBoss Logging.

The logging subsystem is configured using a system of log categories and log handlers. Log categories define what messages to capture, and log handlers define how to deal with those messages (write to disk, send to console etc).

Logging Profiles allow uniquely named sets of logging configuration to be created and assigned to applications independent of any other logging configuration. The configuration of logging profiles is almost identical to the main logging subsystem.

9.2. DEFAULT LOG FILE LOCATIONS

These are the log files that get created for the default logging configurations. The default configuration writes the server log files using periodic log handlers

Table 9.1. Default Log File for a standalone server

Log File	Description
<i>EAP_HOME</i> /standalone/log/server.log	Server Log. Contains all server log messages, including server startup messages.
<i>EAP_HOME</i> /standalone/log/gc.log	Garbage collection log. Contains details of all garbage collection.

Table 9.2. Default Log Files for a managed domain

Log File	Description
<i>EAP_HOME</i> /domain/log/host-controller.log	Host Controller boot log. Contains log messages related to the startup of the host controller.
<i>EAP_HOME</i> /domain/log/process-controller.log	Process controller boot log. Contains log messages related to the startup of the process controller.
<i>EAP_HOME</i> /domain/servers/ <i>SERVERNAME</i> /log/server.log	The server log for the named server. Contains all log messages for that server, including server startup messages.

9.3. JBOSS DATA VIRTUALIZATION LOG CATEGORIES


Within EAP log files, all information specific to JBoss Data Virtualization is prefixed with the `org.teiid` log category identifier.

Logs originating from third-party code and other components not related to JBoss Data Virtualization, including integrated org.jboss components, are prefixed with their own log category identifiers, not `org.teiid`.

The table below summarizes the categories relevant to JBoss Data Virtualization. For a complete listing, see the `standalone-teiid.xml` (or `domain-teiid.xml` for domain mode) file.

Table 9.3. JBoss Data Virtualization Log Categories

Log Category	Description
<code>com.arjuna</code>	Third-party transaction manager. This will include information about all transactions, not just those for JBoss Data Virtualization.
<code>org.teiid</code>	Root category identifier for all JBoss Data Virtualization logs. Note: there are potentially more contexts used under <code>org.teiid</code> than those shown in this table.
<code>org.teiid.PROCESSOR</code>	Query processing logs. Also see <code>org.teiid.PLANNER</code> .
<code>org.teiid.PLANNER</code>	Query planning logs.
<code>org.teiid.SECURITY</code>	Session/Authentication events. Also see <code>org.teiid.AUDIT_LOG</code> .
<code>org.teiid.TRANSPORT</code>	Events related to the socket transport.
<code>org.teiid.RUNTIME</code>	Events related to work management and system start/stop.
<code>org.teiid.CONNECTOR</code>	Connector logs.
<code>org.teiid.BUFFER_MGR</code>	Buffer and storage management logs.
<code>org.teiid.TXN_LOG</code>	Detailed log of all transaction operations.
<code>org.teiid.COMMAND_LOG</code>	Refer to Section 9.4, “Command Logging” .
<code>org.teiid.AUDIT_LOG</code>	Refer to Section 9.5, “Audit Logging” .
<code>org.teiid.ADMIN_API</code>	Admin API logs.
<code>org.teiid.ODBC</code>	ODBC logs.



NOTE

Refer to the JBoss Enterprise Application Platform *Administration and Configuration Guide* for more information about logging.

9.4. COMMAND LOGGING

Command logging captures user commands that have been submitted to JBoss Data Virtualization, query plan commands when query planning is performed, and data source commands that are being executed by connectors.

The user command, "START USER COMMAND", is logged when JBoss Data Virtualization starts working on the query for the first time. This does not include the time the query was waiting in the queue. A corresponding user command, "END USER COMMAND", is logged when the request is complete (that is, when the statement is closed or all the batches are retrieved). There is only one pair of these for every user query.

The query plan command, "PLAN USER COMMAND", is logged when JBoss Data Virtualization finishes the query planning process. There is no corresponding ending log entry.

The data source command, "START DATA SRC COMMAND", is logged when a query is sent to the data source. And a corresponding data source command, "END SRC COMMAND", is logged when the execution is closed (that is, all the rows has been read). There can be one pair for each data source query that has been executed by JBoss Data Virtualization, and there can be number of pairs depending upon your query.

With this information being captured, the overall query execution time can be calculated. Additionally, each source query execution time can be calculated. If the overall query execution time is showing a performance issue, then look at each data source execution time to see where the issue may be.

With this information being captured, the overall query execution time in Teiid can be calculated. Additionally, each source query execution time can be calculated. If the overall query execution time is showing a performance issue, then look at each data source execution time to see where the issue maybe.

To enable command logging to the default log location, simply enable the DETAIL level of logging for the org.teiid.COMMAND_LOG context.

To enable command logging to an alternative file location, configure a separate file appender for the DETAIL logging of the org.teiid.COMMAND_LOG context. An example of this is shown below and can also be found in the standalone-teiid.xml distributed with Teiid.

```
<periodic-rotating-file-handler name="COMMAND_FILE">
  <level name="DEBUG" />
  <formatter>
    <pattern-formatter pattern="%d{HH:mm:ss,SSS} %-5p [%c] (%t)
%s%E%n" />
  </formatter>
  <file relative-to="jboss.server.log.dir" path="command.log" />
  <suffix value=".yyyy-MM-dd" />
</periodic-rotating-file-handler>

<logger category="org.teiid.COMMAND_LOG">
  <level name="DEBUG" />
  <handlers>
    <handler name="COMMAND_FILE" />
  </handlers>
</logger>
```

See *Red Hat JBoss Data Virtualization Development Guide: Server Development* for information on developing a custom logging solution if file based (or any other built-in log4j) logging is not sufficient.

The following is an example of a data source command and what one would look like when printed to the command log:

```
2012-02-22 16:01:53,712 DEBUG [org.teiid.COMMAND_LOG]
(Worker1_QueryProcessorQueue11 START DATA SRC COMMAND: startTime=2012-02-
```



```

22 16:01:53.712
requestID=Ku4/dgtZPYk0.5 sourceCommandID=4 txID=null modelName=DTHCP
translatorName=jdbc-simple sessionID=Ku4/dgtZPYk0
principal=user@teiid-security
sql=HCP_ADDR_XREF.HUB_ADDR_ID, CPN_PROMO_HIST.PROMO_STAT_DT FROM
CPN_PROMO_HIST, HCP_ADDRESS, HCP_ADDR_XREF
WHERE (HCP_ADDRESS.ADDR_ID = CPN_PROMO_HIST.SENT_ADDR_ID) AND
(HCP_ADDRESS.ADDR_ID = HCP_ADDR_XREF.ADDR_ID) AND
(CPN_PROMO_HIST.PROMO_STAT_CD NOT LIKE 'EMAIL%') AND
(CPN_PROMO_HIST.PROMO_STAT_CD <> 'SENT_EM') AND
(CPN_PROMO_HIST.PROMO_STAT_DT > {ts'2010-02-22 16:01:52.928'})

```

Note the following pieces of information:

- **modelName:** this represents the physical model for the data source that the query is being issued.
- **translatorName:** shows type of translator used to communicate to the data source.
- **principal:** shows the user account who submitted the query
- **startTime/endTime:** the time of the action, which is based on the type command being executed.
- **sql:** is the command submitted to the translator for execution, which is NOT necessarily the final sql command submitted to the actual data source. But it does show what the query engine decided to push down.

9.5. AUDIT LOGGING

Audit logging captures important security events, including the enforcement of permissions, and authentication success/failure.

See *Red Hat JBoss Data Virtualization Development Guide: Server Development* for information on developing a custom logging solution if file based (or any other built-in log4j) logging is not sufficient.

9.6. ENABLE AUDIT AND COMMAND LOGGING

You need to enable audit and command logging via the Admin Console. You will also need to grant at least some users the logging role.



NOTE

If using a data base to store the logging, then the Dashboard logging workspace can be used.

The teiidUser requires logging role permissions for this to work (you most likely configured this through the installer.)

9.7. COMMAND LOG OUTPUT

The following is an example of what a data source command looks like when printed to the command log:

```
2012-02-22 16:01:53,712
DEBUG [org.teiid.COMMAND_LOG] (Worker1_QueryProcessorQueue11 START DATA
SRC COMMAND: startTime=2012-02-22 16:01:53.712
requestID=Ku4/dgtZPYk0.5 sourceCommandID=4 txID=null modelName=DTHCP
translatorName=jdbc-simple sessionID=Ku4/dgtZPYk0
principal=user@teiid-security
sql=HCP_ADDR_XREF.HUB_ADDR_ID, CPN_PROMO_HIST.PROMO_STAT_DT FROM
CPN_PROMO_HIST, HCP_ADDRESS, HCP_ADDR_XREF
WHERE (HCP_ADDRESS.ADDR_ID = CPN_PROMO_HIST.SENT_ADDR_ID) AND
(HCP_ADDRESS.ADDR_ID = HCP_ADDR_XREF.ADDR_ID) AND
(CPN_PROMO_HIST.PROMO_STAT_CD NOT LIKE 'EMAIL%') AND
(CPN_PROMO_HIST.PROMO_STAT_CD <> 'SENT_EM') AND
(CPN_PROMO_HIST.PROMO_STAT_DT > {ts'2010-02-22 16:01:52.928'})
```

Note the following pieces of information:

- **modelName** represents the physical model on which the query is being issued.
- **translatorName** shows type of translator used to communicate to the data source.
- **principal** shows the account of the user who submitted the query.
- **startTime/endTime** is the time of the action, which is based on the type of command being executed.
- **sql** is the command submitted to the translator for execution, which is NOT necessarily the final SQL command submitted to the actual data source, but it does show what the query engine decided to push down.

CHAPTER 10. CLUSTERING

10.1. CLUSTERING IN JBOSS DATA VIRTUALIZATION

Clustering may be used to improve the performance of the system through:

Load Balancing

See *Red Hat JBoss Data Virtualization Development Guide: Client Development* for information on load balancing between multiple nodes.

Failover

See *Red Hat JBoss Data Virtualization Development Guide: Client Development* for information on failover with multiple nodes.

Distributed Caching

This is done automatically when clustering is configured.

Event Distribution

Metadata and data modifications will be distributed to all members of a cluster automatically when clustering is configured.

10.2. ENABLE CLUSTERING IN JBOSS DATA VIRTUALIZATION

Ensure JBoss Data Virtualization is installed on each JBoss EAP node and that JBoss EAP has started using either the `standalone-ha.xml` or the `standalone-full-ha.xml` profile before starting the cluster. (These profiles provide high availability, or clustering, as their names imply.)

For more information, see the *Red Hat JBoss Data Virtualization Installation Guide* and Red Hat JBoss Enterprise Application Platform documentation.



NOTE

Artifacts such as VDBs cannot be deployed via file deployment when clustering is used. Use other methods of VDB deployment.

10.3. CONFIGURE LOAD BALANCING

The Teiid JDBC driver does not perform true load-balancing. You can use it to route queries across the host:port combinations defined in the URL but it will not do it based on the load. Instead, you need to use HAProxy.

CHAPTER 11. CACHING

11.1. RESULT SET CACHING

11.1.1. Result Set Caching

JBoss Data Virtualization can cache the results of specific user queries and virtual procedure calls. This caching technique can yield significant performance gains if users of the system submit the same queries or execute the same procedures often.

11.1.2. Result Set Caching Features

Result set caching provides support for the following features:

- User query results can be cached, including XML document model results.
- Virtual procedure results can be cached.
- Results can be scoped automatically as either VDB/user (replicated) or session level.
- Users can configure the number of cache entries and time to live.
- Cache can be cleared via administration tools.

11.1.3. Learn About Caching

To learn about caching in Red Hat JBoss Data Virtualization, please refer to *Development Guide Volume 5: Caching Guide*. This discusses the topic in detail.

PART V. MONITORING AND PERFORMANCE

CHAPTER 12. MONITORING

12.1. MONITORING JBOSS DATA VIRTUALIZATION

JBoss Data Virtualization provides information about its current operational state. This information can be useful in tuning, monitoring, and managing load and throughput. Runtime data can be accessed using administrative tools such as Management Console, AdminShell or Admin API.

12.2. QUERY/SESSION DETAILS

Name	Description
Current Sessions	List current connected sessions
Current Request	List current executing requests
Current Transactions	List current executing transactions
Query Plan	Retrieves the query plan for a specific request

There are administrative options for terminating sessions, queries, and transactions.

12.3. SESSION/QUERY METRICS

Name	Property	Description	Comment
Session Count	sessionCount	Indicates the number of user connections currently active.	To ensure number of sessions are not restricted at peak times, check max-sessions-allowed (default 5000) is set accordingly and review sessions-expiration-timelimit.
Query Count	queryCount	Indicates the number of queries currently active.	
Active Query Plan Count	ENGINE_STATISTIC.active-plans-count	Number of query plans currently being processed.	To ensure maximum throughput, see the performance tuning section on QueryEngine and threading.
Waiting Query Plan Count	ENGINE_STATISTIC.waiting-plans-count	Number of query plans currently waiting.	

Name	Property	Description	Comment
Max Waiting Query Plan Watermark	ENGINE_STATISTIC.max-waitplan-watermark	The maximum number of query plans that have been waiting at one time, since the last time the server started.	
Long Running Queries	longRunningQueries	List current executing queries that have surpassed the query threshold (query-threshold-in-seconds.).	Setup alert to warn when one or more queries are consuming resources for an extended period of time. If running too long, an option is to cancel request or increase threshold.

12.4. BUFFER MANAGER METRICS

Name	Property	Description	Comment
Disk Write Count	ENGINE_STATISTIC.buffermgr-disk-write-count	Disk write count for the buffer manager.	
Disk Read Count	ENGINE_STATISTIC.buffermgr-disk-read-count	Disk read count for the buffer manager.	
Cache Write Count	ENGINE_STATISTIC.buffermgr-cache-write-count	Cache write count for the buffer manager.	
Cache Read Count	ENGINE_STATISTIC.buffermgr-cache-read-count	Cache read count for the buffer manager.	
Disk Space Used (MB)	ENGINE_STATISTIC.buffermgr-diskspace-used-mb	Indicates amount of storage space currently used by buffer files	Setup alert to warn when used buffer space is at an unacceptable level, based on the setting of max-buffer-space
Total memory in use (KB)	ENGINE_STATISTIC.total-memory-inuse-kb	Estimate of the current memory usage in kilobytes.	
Total memory in use by active plans (KB)	ENGINE_STATISTIC.total-memory-inuse-active-plans-kb	Estimate of the current memory usage by active plans in kilobytes	

12.5. CACHE METRICS

Name	Property	Description
Prepared Plan Cache Size	PREPARED_PLAN_CACHE.total-entries	Current number of entries in cache.
Prepared Plan Cache # of Requests	PREPARED_PLAN_CACHE.request-count	Total number of requests made against cache.
Prepared Plan Cache Hit Ratio %	PREPARED_PLAN_CACHE.hit-ratio	Percentage of positive cache hits
ResultSet Cache Size	QUERY_SERVICE_RESULT_SET_CACHE.total-entries	Current number of entries in cache.
ResultSet Cache # of Requests	QUERY_SERVICE_RESULT_SET_CACHE.request-count	Total number of requests made against cache.
ResultSet Cache Hit Ratio %	QUERY_SERVICE_RESULT_SET_CACHE.hit-ratio	Percentage of positive cache hits.

CHAPTER 13. PERFORMANCE TUNING

13.1. USAGE TIPS

- Study the documentation, particularly the Getting Started Guide, to find out what data federation can do for your enterprise.
- Use a standard naming convention for JNDI Data Source names to enable VDBs to be portable across EAP/DV server environments because the JNDI references defined in the VDB do not have to be changed as its being ported from server to server.



NOTE

You can reset the JNDI name by using the AdminConsole, without having to modify the VDB.

b) For consistency, decide on naming conventions so that it does not lead to confusion of the different types. Items to consider include Designer project names, source model names, new model names, function/procedures names, data sources, VDB names, and web service names.

- Standardize on a common JDBC driver version across server instances. That is, if the source is MySQL, all servers should use the same MySQL JDBC driver JAR.

It is preferable to use JDBC type 4 drivers, if you can, for ease of deployment and when doing data previews in Teiid Designer. Otherwise, you can manually change the driver to type 4 or configure the module on the server for the driver, prior to use.

- Include UDF jars with your VDB when packaging it in Teiid Designer as this makes it easier to deploy and does not require you to undertake server module configuration.
- Use VDB Versioning of the format `vdbname.version.vdb`. If you do so VDB updates can be deployed without impacting current processing. It will also be easier for you to visually identify what version of the VDB is utilizing, without having to open it up.
- Pick a suitable deployment strategy:
 - JBoss Operations Network – good for pushing a VDB to a cluster of Teiid instances all at once.
 - Admin shell – a command line tool which can be useful for scripting and automating task, and incorporating into other tools (see example: `DeployAndVerifyVDB.groovy` in the DV kit)
 - Admin-console can be used to deploy the VDB
 - Teiid Designer can deploy the VDB, but Red Hat recommends that you do not use it beyond the development phase
 - Manual copying of artifacts - you are likely to do this only during upgrades or as a means for promoting through your development life cycle.
- With regard to the software development lifecycle, consider the following VDB connection type changes during the deployment life cycle:
 - When a VDB is deployed, use connection type default of `BY_VERSION`

- When the VDB is to be updated on the server with a new version change the current VDB (if exist) connection type from ANY to NONE (to stop allowing connections), so new connections will go to new version and now change the newly deployed VDB connection type from BY_VERSION to ANY, making it the target VDB to be connected to, and positioning it for the next time a new version is to be deployed.
- If rollback is required, change the new (bad) VDB connection type from ANY to NONE and change old VB connection type from NONE to ANY
- When you are in a multi-developer environment, use an SCM (such as github, SVN, or CVS) to manage your artifacts. Red Hat does not recommend any particular source control management system, but does advise you to use one that integrates with JBoss Developer Studio
- Consider using the VDB reuse feature with DV 6 via import-vdb. This allows VDBs to import other VDBs at deploy/startup time. This allows a bigger project to be broken out into separate VDBs that work together at runtime. This will allow developers to work on different VDB's that reference the same underlying models. By partitioning and layering, it can limit the impact of deployment changes.
- Decide on a VDB / Model design pattern.

a) A general design approach is based on building a Data Virtualization Abstraction layer (also referred to as Enterprise or Federated Data Layer) The Abstraction Layer is what client applications use. Its referred to as the abstraction layer because the underlying data sources are abstracted away from client use. A recommended approach to building the abstraction layer is to start by creating a Virtual Base Layer (VBL) between the 'Data Source Layer' and the 'Abstraction Layer'. This decouples the physical data from the canonical business model, allowing flexibility to change the data layer without affecting the client and providing a level of isolation on top of the physical sources. This is a one-to-one mapping of each physical source to a virtual model. These VBL components can then be used as building blocks for higher level services; all of the transformations built on top of them will not need to be changed should the need arise to accommodate a change in the source(s)

b) In laying out your design pattern, consider using VDB reuse (multiple VDBs) such as VDBs at the abstraction layer and VDBs at the physical layer

Performance Tips

- Consider your naming conventions. The naming convention you use can reduce the size of messages.
- Remember to undeploy any unused old VDBs.

JBDS Teiid Plug-In Tips

- When you are developing in a collaborative environment, developers need to use the same project names. There can be multiple projects (for example, by application) and there can be sharing of models across projects when creating a VDB. However, in the designer, the references from a VDB to a model includes the project name, so using a different project name will result in the VDB's models not being resolved correctly.

Establish a standard naming convention and configuration for Connection Profiles across eclipse/JBDS environments Otherwise, when a model is imported and the connection profile does not match, the connection profile will need to be reset for the source model before data previewing or VDB deployment can be done.

- When modelling, and you go to build a View model, create children of type "Base Table". Do not use "View Table". There is no functional difference. The problem comes in later when you want to create PK/FK/Indexes on the table. The "View Table" does not allow you to do this.
 - (a) If you are going to be doing updates, it is recommended to use XA data sources. Therefore, do not use Designer to create the data sources. This is due to a limitation in DTP, which will create a 'local' data source, not an XA data source. To create XA data source, look in JBoss AS "doc" directory for example templates, or use the "admin-console" to create the XA data sources

Tips for Custom Translators and Connectors

- Use the Teiid translator and/or connector archetypes.

When used, they will create the project template, with resources, to immediately begin adding your custom code and compiling to package for deployment.

Using JBDS, create a new Maven project by creating the project from the archetype. In the wizard, select the "JBoss Nexus Archetypes" catalog, and filter on "teiid".

13.2. MEMORY MANAGEMENT CONSIDERATIONS

Here is some information about settings related to memory management.

The JBoss Data Virtualization engine uses batching to reduce the number of memory rows processed at a given time. The batch sizes may be adjusted to larger values as more clients will be accessing the server simultaneously.

buffer-service-max-reserve-kb

Default is -1. This setting determines the total size in kilobytes of batches that can be held by the buffer manager in memory. This number does not account for persistent batches held by soft (such as index pages) or weak references. The default value of -1 will automatically calculate a typical maximum based on the maximum heap available to the VM. The calculated value assumes a 64bit architecture and will limit buffer usage to 50% of the first gigabyte of memory beyond the first 300 megabytes (which are assumed for use by JBoss EAP) and 75% of the memory beyond that.

For example, with default settings and an 8GB VM size, then **buffer-service-max-reserve-kb** will at a maximum use: $((1024-300) * 0.5) + (7 * 1024 * 0.75) = 5738 \text{ MB or } 5875712 \text{ KB}$.

The buffer manager automatically triggers the use of a canonical value cache if enabled when more than 25% of the **buffer-service-max-reserve-kb** is in use. This can dramatically cut memory usage in situations where similar value sets are being read by JBoss Data Virtualization, but does introduce a lookup cost. If you are processing small or highly similar datasets using JBoss Data Virtualization, and wish to conserve memory, you should consider enabling value caching.



NOTE

Memory consumption can be significantly more or less than the nominal target depending on actual column values and whether value caching is enabled. Large non built-in type objects can exceed their default size estimate. If out of memory errors occur, lower the **buffer-service-max-reserve-kb** value. Also note that source LOB values are held by memory references that are not cleared when a batch is persisted. With heavy LOB usage you should ensure that buffers of other memory associated with LOB references are appropriately sized.

buffer-service-max-processing-kb

Default is -1. This setting determines the total size in kilobytes of batches that can be guaranteed for use by one active plan and may be in addition to the memory held based on **buffer-service-max-reserve-kb**. Typical minimum memory required by JBoss Data Virtualization when all the active plans are active is $\#(\text{active-plans}) * \text{buffer-service-max-processing-kb}$. The default value of -1 will automatically calculate a typical maximum based upon the maximum heap available to the VM and max active plans. The calculated value assumes a 64bit architecture and will limit processing batch usage to 10% of memory beyond the first 300 megabytes (which are assumed for use by JBoss EAP).

For example, with default settings including 20 active-plans and an 8GB VM size, then **buffer-service-max-processing-kb** will be: $((1024-300) * 0.1) + (7 * 1024 * 0.1))/20 = 789.2 \text{ MB}/20 = 39.46 \text{ MB}$ or 40407 KB per plan. This implies a range between 0 and 789 MB that may be reserved with roughly 40 MB per plan.

In systems where large intermediate results are expected you can consider increasing **buffer-service-max-processing-kb** and decreasing **buffer-service-max-reserve-kb** so that each request has access to an effectively smaller buffer space.

buffer-service-max-file-size

Default is 2GB. Each intermediate result buffer, temporary LOB, and temporary table is stored in its own set of buffer files, where an individual file is limited to **buffer-service-max-file-size** megabytes. Consider increasing the storage space available to all such files using **buffer-service-max-buffer-space** if your installation makes use of internal materialization, makes heavy use of SQL/XML, or processes large row counts.

TEIID_AUDIT_LOGGING and TEIID_COMMAND_LOGGING

The logging queue-length for **TEIID_AUDIT_LOGGING** and **TEIID_COMMAND_LOGGING** should be increased. As a minimum, you must set it to the expected concurrency level for that node. In other words, if you expect 1,000 users to use the system concurrently then set each of the queue-lengths to 1k and then tune it from there.

Limitations

It is also important to keep in mind that Teiid has memory and other hard limits which breaks down along several lines in terms of # of storage objects tracked, disk storage, streaming data size/row limits, and so forth.

The buffer manager has a max addressable space of 16 terabytes - but due to fragmentation you'd expect that the max usable would be less. This is the maximum amount of storage available to Teiid for all temporary lobbs, internal tables, intermediate results, etc.

The max size of an object (batch or table page) that can be serialized by the buffer manager is 32 GB - but you should not get near that (the default limit is 8 MB). A batch is set or rows that are flowing through Teiid engine. Teiid temporary tables (also used for internal materialization) can only support $2^{31}-1$ rows per table.

However handling a source that has tera/petabytes of data doesn't by itself impact Teiid in any way. What matters is the processing operations that are being performed and/or how much of that data do we need to store on a temporary basis in Teiid. With a simple forward-only query, as long as the result row count is less than 2^{31} , Teiid be perfectly happy to return a petabyte of data.

Other Considerations for Sizing

Each batch/table page requires an in memory cache entry of approximately ~ 128 bytes - thus the total tracked max batches are limited by the heap and is also why we recommend to increase the processing batch size on larger memory or scenarios making use of large internal materializations. The actual batch/table itself is managed by buffer manager, which has layered memory buffer structure with spill over facility to disk.

Using internal materialization is based on the buffermanager. Buffermanager settings may need to be updated based upon the desired amount of internal materialization performed by deployed VDBs.

13.3. SCALABILITY CONSIDERATIONS

Although, you can find information about all JBoss Data Virtualization settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about those settings related to scalability.

buffer-service-processor-batch-size

Default is 256. This property specifies the maximum row count of a batch sent internally within the query processor. Additional considerations are needed if extremely large VM sizes and or datasets are being used. JBoss Data Virtualization has a non-negligible amount of overhead per batch/table page on the order of 100-200 bytes. Depending on the data types involved, each full batch/table page will represent a variable number of rows (a power of two multiple above or below the processor batch size).


If you are working with extremely large datasets and you run into memory issues, consider increasing the **buffer-service-processor-batch-size** property to force the allocation of larger batches and table pages.

buffer-service-max-storage-object-size

Default is 8288608 or 8MB. This value is the maximum size of a buffered managed object in bytes and represents the individual batch page size.

If **buffer-service-processor-batch-size** is increased or you are dealing with extremely wide result sets, then the default setting of 8MB for the **buffer-service-max-storage-object-size** may be too low. The inline LOBS also contribute to this size if the batch contains them. The sizing for **buffer-service-max-storage-object-size** is in terms of serialized size, which will be much closer to the raw data size than the Java memory footprint estimation used for **buffer-service-max-reserve-kb**.

buffer-service-max-storage-object-size should not be set too large relative to **buffer-service-memory-buffer-space** since it will reduce the performance of the memory buffer. The memory buffer supports only 1 concurrent writer for each **buffer-service-max-storage-object-size** of the **buffer-service-memory-buffer-space**.




NOTE

JBoss Data Virtualization temporary tables (also used for internal materialization) can only support $2^{31}-1$ rows per table.

buffer-service-memory-buffer-space

Default is -1. This controls the amount of on or off heap memory allocated as byte buffers for use by the JBoss Data Virtualization buffer manager. This setting defaults to -1, which automatically determines a setting based upon whether it is on or off heap and the value for **buffer-service-**

max-reserve-kb.



NOTE

When left at the default setting the calculated memory buffer space will be approximately one quarter of the **buffer-service-max-reserve-kb** setting. If the memory buffer is off heap and the **buffer-service-max-reserve-kb** setting is automatically calculated, the memory buffer space will be subtracted out of the effective **buffer-service-max-reserve-kb**.

buffer-service-memory-buffer-off-heap

Default is false. Determines whether to take advantage of the buffer manager memory buffer to access system memory without allocating it to the heap. Setting **buffer-service-memory-buffer-off-heap** to **true** will allocate the JBoss Data Virtualization memory buffer off heap. Depending on whether your installation is dedicated to JBoss Data Virtualization and the amount of system memory available, this may be preferred over on-heap allocation.

The primary benefit of off-heap memory is additional memory usage for JBoss Data Virtualization without additional garbage collection tuning. This becomes especially important in situations where more than 32GB of memory is desired for the JVM. Note that when using off-heap allocation, the memory must still be available to the java process and that setting the value of **buffer-service-memory-buffer-space** too high may cause the JVM to swap rather than reside in memory. With large off-heap buffer sizes (greater than several gigabytes) you may also need to adjust JVM settings.

For Sun JVMs the relevant JVM settings are **MaxDirectMemorySize** and **UseLargePages**. For example adding:

```
-XX:MaxDirectMemorySize=12g -XX:+UseLargePages
```

to the JVM process arguments would allow for an effective allocation of an (approximately) 11GB JBoss Data Virtualization memory buffer (the **buffer-service-memory-buffer-space** property) accounting for any additional direct memory that may be needed by JBoss EAP or applications running within it.

13.4. DISK USAGE CONSIDERATIONS

Although, you can find information about all JBoss Data Virtualization settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about those settings related to disk usage.

buffer-service-max-buffer-space

Default is 51200. For table page and result batches, the buffer manager will limit the number of files that are dedicated to a particular storage size. However, creation of Large Object (LOB) values (for example through SQL/XML) will typically create one file per LOB once the LOB exceeds the allowable in memory size of 8KB. In heavy usage scenarios, consider pointing the buffer directory on a partition that is routinely defragmented. By default, JBoss Data Virtualization will use up to 50GB of disk space. This is tracked in terms of the number of bytes written by JBoss Data Virtualization. For large data sets, you may need to increase the **buffer-service-max-buffer-space** property.

13.5. THREADING CONSIDERATIONS

Although, you can find information about all JBoss Data Virtualization settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about those settings related to threading.

max-threads

Default is 64. The query engine has several settings that determine its thread utilization. `max-threads` sets the total number of threads available in the process pool for query engine work (such as processing plans, transaction control operations, and processing source queries).

You should consider increasing the maximum threads on systems with a large number of available processors and/or when it is necessary to issue non-transactional queries involving a large number of concurrent source requests.

max-active-plans

Default is 20. This value should always be smaller than `max-threads`. By default, `thread-count-for-source-concurrency` is calculated by $(\text{max-threads} / \text{max_active_plans}) * 2$ to determine the threads available for processing concurrent source requests for each user query. Increasing the `max-active-plans` should be considered for workloads with a high number of long running queries and/or systems with a large number of available processors. If memory issues arise from increasing the `max-threads` and `max-active-plans`, then consider decreasing the amount of heap held by the buffer manager or decreasing the `processor-batch-size` to limit the base number of memory rows consumed by each plan.

Increasing `max-active-plans` should be considered for workloads with a high number of long running queries and/or systems with a large number of available processors. If memory issues arise from increasing `max-threads` and `max-active-plans`, then consider decreasing `buffer-service-processor-batch-size` to limit the base number of memory rows consumed by each plan.

thread-count-for-source-concurrency

Default is 0. This value should always be smaller than `max-threads`. This property sets the number of concurrently executing source queries per user request. 0 indicates to use the default calculated value based on $2 * (\text{max-threads} / \text{max-active-plans})$. Setting this to 1 forces serial execution of all source queries by the processing thread. Any number greater than 1 limits the maximum number of concurrently executing source requests accordingly.

Using the defaults, each user request would be allowed 6 concurrently executing source queries. If the default calculated value is not applicable to your workload, for example, if you have queries that generate more concurrent long running source queries, you should adjust this value.

Also see [Section 13.7, “Transport Considerations”](#) for `max-socket-threads`.

13.6. CACHING CONSIDERATIONS

Although, you can find information about all JBoss Data Virtualization settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about those settings related to caching.

JBoss Data Virtualization settings regarding cache tuning are divided amongst:

- Resultset Cache Tuning

- Prepared Plan Cache Tuning

Cache statistics can be obtained through the Management Console or AdminShell. The statistics can be used to help tune cache parameters and ensure a hit ratio.

Plans are currently fully held in memory and may have a significant memory footprint. When making extensive use of prepared statements and/or virtual procedures, the size of the plan cache may be increased proportionally to number of gigabytes intended for use by JBoss Data Virtualization.

While the result cache parameters control the cache result entries (such as max number, and eviction), the result batches themselves are accessed through the buffer manager. If the size of the result cache is increased, you may need to tune the buffer manager configuration to ensure there is enough buffer space.

Result set and prepared plan caches have their entries invalidated by data and metadata events. By default, these events are captured by running commands through JBoss Data Virtualization (see the *Red Hat JBoss Data Virtualization Developer Guide* for further customization). JBoss Data Virtualization stores compiled forms of update plans or trigger actions with the prepared plan so that if metadata changes, the changes may take effect immediately.

The default `resultset-cache-max-staleness` for resultset caching is 60 seconds to improve efficiency with rapidly changing sources. Consider decreasing this value to make the resultset cache more consistent with the underlying data. Even with a setting of 0, full transactional consistency is not guaranteed.



WARNING

Disabling or constraining these caches will lead to poor performance.

13.7. TRANSPORT CONSIDERATIONS

Although, you can find information about all JBoss Data Virtualization settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about those settings related to transports.

JBoss Data Virtualization provides three transports by default: `odbc`, `jdbc` and `embedded`. Transport settings (such as those listed below) are configured for each.

`max-socket-threads`

Default is 0. Determines the maximum number of threads dedicated to the initial request processing. Zero indicates to use the system default of maximum available processors. Socket threads handle NIO non-blocking IO operations as well as directly servicing any operation that can run without blocking. For longer running operations, the socket threads queue works with the query engine. (The query engine has two properties that determine its thread utilization: `max-threads` and `max-active-plans`.)

All JDBC/ODBC socket operations are non-blocking, so setting the number of `max-socket-threads` higher than the maximum effective parallelism of the machine should not result in greater performance.

ODBC queries are executed synchronously from the socket thread. Simultaneous long-running queries may exhaust the available threads. If you expect a higher concurrent load of long-running queries, consider increasing the default `max-socket-threads` (for example, to 15).

`input-buffer-size`

Default is 0 which will use the system default. Before adjusting `input-buffer-size` for any of the transports, keep in mind that each client will create a new socket connection. Increasing this value should only be done if the number of clients is constrained.

`output-buffer-size`

Default is 0 which will use the system default. Before adjusting `output-buffer-size` for any of the transports, keep in mind that each client will create a new socket connection. Increasing this value should only be done if the number of clients is constrained.

JDBC clients may need to adjust low-level transport values, in addition to SSL client connection properties via a `teiid-client-settings.properties` file placed in the client application's classpath. (An example file can be found within the `EAP_HOME/modules/system/layers/base/org/jboss/teiid/client/main/teiid-client-VERSION.jar` file.)



NOTE

Typical installations will not need to have any of these settings adjusted.

13.8. LARGE OBJECTS (LOBs)

Large Objects (LOBs) consist of data. The three main large object runtime data types used by JBoss are:

1. Binary (BLOB)
 - Contains multimedia objects such as images and audio.
2. Character (CLOB)
 - Contains ASCII characters.
3. Extensible Markup Language (XML)
 - Contains textual data.

LOBs and JBoss

The JBoss Data Services Connector API returns a reference to the LOB if allowed by the JBoss Data Services server. The JBoss Data Services server or JDBC driver can then access the data via a stream rather than retrieving the data all at once. This is useful for several reasons:

- Reduces memory usage when returning the result set to the user.
- Improves performance by passing less data in the result set.
- Enables access to LOBs when required rather than assuming that users will always use the LOB data.

- Enables handling of arbitrarily large data values within a fixed JBoss Data Services memory usage.

These benefits are achieved if the Connector itself does not materialize an entire LOB all at once. For example, the JDBC API supports a streaming interface for BLOB and CLOB data.

Source LOB values are typically accessed by reference, rather than having the value copied to a temporary location. Care must be taken to ensure that source LOBs are returned in a memory-safe manner.

LOBs are broken into pieces when being created and streamed. The size of each piece when fetched by the client can be configured.

Cached lobs will be copied rather than relying on the reference to the source lob.

Temporary lobs created by Teiid will be cleaned up when the result set or statement is closed. To rely on implicit garbage collection based cleanup instead of statement close, the Teiid session variable `clean_lob_onclose` can be set to false (by issuing the query "SELECT teiid_session_set('clean_lob_onclose', false)" - which can be done for example via the new connection sql in the datasource definition). This can be used for local client scenarios that relied on the implicit behavior, such as Designer generated REST VDBs.

13.9. LOB CONSIDERATIONS

Although, you can find information about all JBoss Data Virtualization settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about those settings related to large objects (LOBs).

lob-chunk-size-in-kb

LOBs and XML documents are streamed from the JBoss Data Virtualization server to the JDBC API. Normally, these values are not materialized in the server memory, avoiding potential out-of-memory issues. When using style sheets, or XQuery, whole XML documents must be materialized on the server. Even when using the XMLQuery or XMLTable functions and document projection is applied, memory issues may occur for large documents.

LOBs are broken into pieces when being created and streamed. The maximum size of each piece when fetched by the client can be configured with the **lob-chunk-size-in-kb** property.

The default value is 100. When dealing with extremely large LOBs, you may consider increasing **lob-chunk-size-in-kb** to decrease the amount of round-trips to stream the result. Setting the value too high may cause the server or client to have memory issues.

Source LOB values are typically accessed by reference, rather than having the value copied to a temporary location. Thus care must be taken to ensure that source LOBs are returned in a memory-safe manner. This caution is more for the source driver vendors not to consume VM memory for LOBs.

13.10. OTHER PERFORMANCE TUNING CONSIDERATIONS

Although, you can find information about all JBoss Data Services settings using the Management CLI (see [Section 14.1, “JBoss Data Virtualization Settings”](#)), this section provides some additional information about the **max-source-rows** setting.

max-source-rows

When using JBoss Data Services in a development environment, you may consider setting `max-source-rows` to a small value (for example, 10000) to prevent large amounts of data from being pulled from sources. Leaving the `exception-on-max-source-rows` property set to `true` will alert the developer through an exception that an attempt was made to retrieve more than the specified number of rows.

PART VI. REFERENCE

CHAPTER 14. GENERAL CONFIGURATION

14.1. JBOSS DATA VIRTUALIZATION SETTINGS

The following types of JBoss Data Virtualization settings are available for viewing and modification:

- Buffer service settings
- Cache settings (including result set and prepared plan cache settings)
- Runtime engine deployer settings
- Authorization validator and policy decider settings
- Transport and SSL settings
- Translator settings

To view all of the available settings for JBoss Data Virtualization, run the following command within the Management CLI:

```
/subsystem=teiid:read-resource-description
```



NOTE

To look further into translator and transport (including SSL) settings see [Section 14.4, “Managing Transport and SSL Settings Using Management CLI”](#) and [Section 14.5, “Managing Translator Settings Using Management CLI”](#).

14.2. VIEWING JBOSS DATA VIRTUALIZATION SETTINGS USING MANAGEMENT CLI

To view all of the current settings for JBoss Data Virtualization, run the following command within the Management CLI:


```
/subsystem=teiid:read-resource
```

To view a particular JBoss Data Virtualization setting, run the following command within the Management CLI:

```
/subsystem=teiid:read-attribute(name=SETTING_NAME)
```

For example:

```
/subsystem=teiid:read-attribute(name=max-active-plans)
```



NOTE

To view current translator and transport (including SSL) settings refer to [Section 14.4, “Managing Transport and SSL Settings Using Management CLI”](#) and [Section 14.5, “Managing Translator Settings Using Management CLI”](#).

14.3. CHANGING JBOSS DATA VIRTUALIZATION SETTINGS USING MANAGEMENT CLI

To edit a particular JBoss Data Virtualization setting, run the following command within the Management CLI:

```
/subsystem=teiid:write-attribute(name=SETTING_NAME, value=VALUE)
```



NOTE

To change translator and transport (including SSL) settings refer to [Section 14.4, “Managing Transport and SSL Settings Using Management CLI”](#) and [Section 14.5, “Managing Translator Settings Using Management CLI”](#).

For example:

```
/subsystem=teiid:write-attribute(name=max-active-plans, value=50)
```

For more information about the Management CLI, refer to the Red Hat JBoss Enterprise Application Platform *Administration and Configuration Guide*



NOTE

After modifying any of these settings, you will be prompted to reload the server. Changes to these settings will not take effect until the server is restarted. You can reload the server by running the `reload` command from within the Management CLI. After the server reloads, you may have to reconnect by running the `connect` command to continue working within the Management CLI.

14.4. MANAGING TRANSPORT AND SSL SETTINGS USING MANAGEMENT CLI

To manage JBoss Data Virtualization transport settings, you can use the same commands as those used for the base JBoss Data Virtualization settings, specifying a particular transport in the command. For example:

```
/subsystem=teiid/transport=TRANSPORT_NAME:read-resource
```

Available transport names are listed under `transport` when you run the following command to output current JBoss Data Virtualization settings:

```
/subsystem=teiid:read-resource
```

14.5. MANAGING TRANSLATOR SETTINGS USING MANAGEMENT CLI

To manage JBoss Data Virtualization translator settings, you can use the same commands as those used for the base JBoss Data Virtualization settings, specifying a particular translator in the command. For example:

```
/subsystem=teiid/translator=TRANSLATOR_NAME:read-resource
```

Available translator names are listed under `translator` when you run the following command to output current JBoss Data Virtualization settings:

```
/subsystem=teiid:read-resource
```

14.6. TRANSPORT SECURITY AUTHENTICATION MODES

The following authentication modes are available:

anonymous

No certificates are exchanged. Settings are not needed for the keystore and truststore properties. The client must have `org.teiid.ssl.allowAnon` set to true (the default) to connect to an anonymous server. Communications are encrypted using the `TLS_DH_anon_WITH_AES_128_CBC_SHA` SSL cipher suite. This is suitable for most secure intranets.

1-way

Authenticates the server to the client. The server presents a certificate which is signed by the private key stored in the server's keystore. The server's corresponding public key must be in the client's truststore.

2-way

Mutual client and server authentication. The server presents a certificate which is signed by the private key stored in the server's keystore. The server's corresponding public key must be in the client's truststore. Additionally, the client presents a certificate signed by its private key stored in the client's keystore. The client's corresponding public key must be in the server's truststore.



NOTE

You can use `keytool` to generate encryption keys; however, you should first consider your local requirements for managing public key cryptography.

14.7. MANAGING CORE CONFIGURATION USING JBOSS MANAGEMENT CONSOLE

1. Log in to the JBoss Management Console.
2. Select the **Profile** tab.
3. From the navigation tree, select **Subsystems** → **Teiid**.
4. Select one of the following options:
 - Query Engine - From here you can view and configure core query engine properties.
 - Translators - From here you can view, add and remove translators.
 - Transports - From here you can view, add and remove transports.
5. View and modify configuration as necessary.

**NOTE**

If you need help at any stage, select the **Need Help?** link.

**WARNING**

Some properties require you to restart the server before they will take effect.

14.8. PORTS USED BY RED HAT JBOSS DATA VIRTUALIZATION

Red Hat JBoss Data Virtualization inherits the ports used by JBoss Enterprise Application Platform. You can find a full list here: https://access.redhat.com/documentation/en-US/JBoss_Enterprise_Application_Platform/6/html/Administration_and_Configuration_Guide/Network_

In addition, to these, ports 31000 and 35432 must also be kept open.

14.9. CHANGE THE DEFAULT JDBC PORT USING MANAGEMENT CONSOLE

1. Login to the Management Console.
2. Navigate to the **Socket Binding** panel in the Management Console
 - a. **Standalone Mode**
Select the **Profile** tab from the top-right of the console.
 - Domain Mode**
 - i. Select the **Profiles** tab from the top-right of the console.
 - ii. Select the appropriate profile from the drop-down box in the top left.
 - iii. Expand the **Subsystems** menu on the left of the console.
 - b. Select **General Configuration** → **Socket Binding** from the menu on the left of the console.
3. **Modify the port number**
 - a. Select the `teiid-jdbc` configuration.
 - b. Select the **Edit** button.
 - c. Set the **Port** to the new port number.
 - d. Select **Save**.

14.10. SYSTEM PROPERTIES


Some behavior can be configured via system properties, rather than configuration files. A typical place to set system properties for JBoss Data Virtualization is in the `EAP_HOME/bin/MODE.conf`. A property setting has the format `-Dproperty=value`.

Table 14.1. System Properties

Setting	Description	Default Value
<code>org.teiid.allowNaNInfinity</code>	Set to true to allow numeric functions to return NaN (Not A Number) and +-Infinity. Note that these values are not covered by the SQL specification.	Defaults to false.
<code>org.teiid.useValueCache</code>	Set to true to enable the canonical value cache. Value caching is used dynamically when buffer memory is running low to reuse identical values, reducing the memory consumed by JBoss Data Virtualization. However, there is a computation cost associated with the cache lookup, so enabling this setting is not appropriate for installations handling large volumes of dissimilar data.	Defaults to false.
<code>org.teiid.ansiQuotedIdentifiers</code>	Set to false to emulate prior behavior of treating double quoted values without leading identifier parts as string literals, which is not expected by the SQL specification.	Defaults to true.
<code>org.teiid.subqueryUnnestDefault</code>	Set to true to aggressively unnest subquery IN and EXISTS predicates. If possible, the predicate will be unnested to a traditional join and will be eligible for dependent join planning. If a traditional join is not possible (such as with NOT IN) a merge join version of the semijoin or antijoin will be considered based upon the costing information available.	Defaults to false.
<code>org.teiid.ODBCPacketSize</code>	Target size in bytes of the ODBC results buffer. This is not a hard maximum, LOBS and wide rows may use larger buffers.	Defaults to 307200.

Setting	Description	Default Value
<code>org.teiid.decimalAsDouble</code>	Set to true to parse exact fixed point literals (for example, 1.0) as double values rather than as decimal/BigDecimal values and to return a double value from the AVG function for integral values in the same way as previous versions.	Defaults to false.
<code>org.teiid.comparableLobs</code>	Set to true to allow BLOB and CLOB column values to be comparable in JBoss Data Virtualization. Source type metadata will determine if the comparison can be pushed down.	Defaults to false.
<code>org.teiid.comparableObject</code>	Set to true to allow object column values to be comparable in JBoss Data Virtualization. Source type metadata will determine if the comparison can be pushed down. The object instances are expected to correctly implement <code>java.lang.Comparable.compareTo</code> . If the instance object is not <code>Comparable</code> , then <code>ClassCastException</code> s may be thrown.	Defaults to false.
<code>org.teiid.padSpace</code>	Set to true to compare strings as if PAD SPACE collation is being used. This means strings will be right padded to the same length for comparison. If this property is set, it is not necessary to use the <code>trimStrings</code> translator option.	Defaults to false.
<code>org.teiid.collationLocale</code>	Set to a Java locale string <code>language[_country][_variant]]</code> , where language, country, and variant are two letter codes - see java.util.Locale for more on valid codes. Note that even if <code>org.teiid.comparableLobs</code> is set, CLOB values will not be compared using the locale collator.	Not set by default, which means that Java's natural (UTF-16) string comparison will be used.

Setting	Description	Default Value
<code>org.teiid.clientVdbLoadTimeoutMillis</code>	The default amount of time a client (currently only local clients) will wait to make a connection to an active VDB before throwing an exception. Clients may override this setting via the <code>waitForLoad</code> connection property.	Defaults to 5 minutes.
<code>org.teiid.enDateNames</code>	Set to true to use English month and day names for the system function <code>dayName</code> and <code>monthName</code> , rather than returning names from the Java default locale. Prior to this release, <code>dayName</code> and <code>monthName</code> always returned English names.	Defaults to false.
<code>org.teiid.pushdownDefaultNullOrder</code>	Set to true to mimic prior release behavior of pushing the default null order of nulls low if the source has a different default null order and supports explicit null ordering.	Defaults to false.
<code>org.teiid.implicitMultiSourceJoin</code>	Set to false to disable prior release behavior of implicitly partitioning joins between multi-source tables. When set to false an explicit predicate such as <code>tbl1.source_name = tbl2.source_name</code> is required to partition the results of the join.	Defaults to true.
<code>org.teiid.joinPrefetchBatches</code>	Sets the number of batches that can be pre-fetched/buffered for join processing that may otherwise be streamed from the source. Consider increasing when you have slow <code>ALREADY_SORTED</code> sources participating in a non-dependent join that return significantly more batches. Note however that increasing this value can lead to an increase in disk utilization to perform the buffering.	Defaults to 10.

Setting	Description	Default Value
org.teiid.maxStringLength	<p>Sets the nominal maximum length of strings in JBoss Data Virtualization. Most operations will truncate strings that are larger than this value. Setting this value can also adjust the maximum size of LOB bytes held in memory. Note that sources may not appropriately handle string values that are larger than what the source supports.</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;">  <p>WARNING</p> <p>Strings are held entirely in memory. Do not set this value too high as you may experience out of memory errors.</p> </div>	Defaults to 4000.
org.teiid.calendarTimestampDiff	Set to false to use the timestampdiff behaviour from previous versions. Note that using the old behavior can result in differing results between pushed and non-pushed versions of timestampdiff for intervals greater than seconds. This is because sources use the date part and not approximate interval differences.	Defaults to true

The following properties are provided for PostgreSQL compatibility.

Table 14.2. System Properties for PostgreSQL Compatibility

Setting	Description	Default Value
---------	-------------	---------------

Setting	Description	Default Value
org.teiid.iso8601Week	Set to true to use ISO 8601 rules for week calculations regardless of the locale. When set to true, the dayOfWeek function will begin with 1 for MONDAY rather than SUNDAY, and the week function will require that week 1 of a year contains the year's first Thursday.	Defaults to false.
org.teiid.backslashDefaultMatchEscape	Set to true to use '\' as the default escape character for LIKE and SIMILAR TO predicates when no escape is specified. Otherwise JBoss Data Virtualization assumes the SQL specification compliant behavior of treating each non-wildcard character as an exact match character.	Defaults to false.

**NOTE**

These properties for PostgreSQL compatibility affect JBoss Data Virtualization globally, not just through the ODBC transport.

14.11. TEIID MANAGEMENT CLI

The AS CLI is a command line based administrative and monitoring tool for Teiid. AdminShell provides a binding into the Groovy scripting language and higher level methods that are often needed when interacting with Teiid. It is still useful to know the underlying CLI commands in many circumstances. The below is a series useful CLI commands for administering a Teiid Server.

VDB Operations

```

deploy adminapi-test-vdb.xml
undeploy adminapi-test-vdb.xml

/subsystem=teiid:restart-vdb(vdb-name=AdminAPITestVDB, vdb-version=1,
model-names=TestModel)

/subsystem=teiid:list-vdbs()
/subsystem=teiid:get-vdb(vdb-name=AdminAPITestVDB, vdb-version=1)
/subsystem=teiid:change-vdb-connection-type(vdb-name=AdminAPITestVDB, vdb-
version=1, connection-type=ANY)

/subsystem=teiid:add-data-role(vdb-name=AdminAPITestVDB, vdb-version=1,
data-role=TestDataRole, mapped-role=test)
/subsystem=teiid:remove-data-role(vdb-name=AdminAPITestVDB, vdb-version=1,
data-role=TestDataRole, mapped-role=test)

```

Source Operations

```
/subsystem=teiid:add-source(vdb-name=AdminAPITestVDB, vdb-version=1,
source-name=text-connector-test, translator-name=file, model-
name=TestModel, ds-name=java:/test-file)
/subsystem=teiid:remove-source(vdb-name=AdminAPITestVDB, vdb-version=1,
source-name=text-connector-test, model-name=TestModel)
/subsystem=teiid:update-source(vdb-name=AdminAPITestVDB, vdb-version=1,
source-name=text-connector-test, translator-name=file, ds-
name=java:/marketdata-file)
```

Translator Operations

```
/subsystem=teiid:list-translators()
/subsystem=teiid:get-translator(translator-name=file)
/subsystem=teiid:read-translator-properties(translator-
name=file, type=OVERRIDE)
/subsystem=teiid:read-rar-description(rar-name=file)
```

Runtime Operations

```
/subsystem=teiid:workerpool-statistics()

/subsystem=teiid:cache-types()
/subsystem=teiid:clear-cache(cache-type=PREPARED_PLAN_CACHE)
/subsystem=teiid:clear-cache(cache-type=QUERY_SERVICE_RESULT_SET_CACHE)
/subsystem=teiid:clear-cache(cache-type=PREPARED_PLAN_CACHE, vdb-
name=AdminAPITestVDB, vdb-version=1)
/subsystem=teiid:clear-cache(cache-type=QUERY_SERVICE_RESULT_SET_CACHE,
vdb-name=AdminAPITestVDB, vdb-version=1)
/subsystem=teiid:cache-statistics(cache-type=PREPARED_PLAN_CACHE)
/subsystem=teiid:cache-statistics(cache-
type=QUERY_SERVICE_RESULT_SET_CACHE)

/subsystem=teiid:engine-statistics()

/subsystem=teiid:list-sessions()
/subsystem=teiid:terminate-session(session=sessionid)

/subsystem=teiid:list-requests()
/subsystem=teiid:cancel-request(session=sessionid, execution-id=1)
/subsystem=teiid:list-requests-per-session(session=sessionid)
/subsystem=teiid:list-transactions()

/subsystem=teiid:mark-datasource-available(ds-name=java:/accounts-ds)

/subsystem=teiid:get-query-plan(session=sessionid, execution-id=1)
```

CHAPTER 15. DIRECTORY STRUCTURE

15.1. DIRECTORY STRUCTURE

The following shows the contents of the JBoss Data Virtualization deployment within an EAP instance.

```

/bin
  /scripts
/docs
  /teiid
    /datsources
    /schema
    /examples
/domain
  /configuration
    application-users.properties
    application-roles.properties
/dataVirtualization
/modules
  /system
    /layers
      /dv
        /dv modules
      /base
        /eap modules
/standalone
  /configuration
    standalone.xml
    application-users.properties
    application-roles.properties

```

bin/scripts

This directory contains installation and utility CLI scripts.

docs/teiid

This directory contains documents, examples, sample data source XML fragments and schema files.

standalone/configuration

- `standalone.xml` is the master configuration file for JBoss Data Virtualization. This file manages configuration for the JBoss Data Virtualization subsystem in addition to the standard JBoss EAP web profile subsystems.
- `application-users.properties` and `application-roles.properties` define the allowed users and their defined roles using the default security domain. Edit these files to add users. If you want to use a different security domain, change the security domain details in the main configuration file.

domain/configuration

- `application-users.properties` and `application-roles.properties` define the allowed users and their defined roles in JBoss Data Virtualization using the default security domain. Edit these files to add users. If you want to use a different security domain, change

the security domain details in main configuration file.

/modules/system/layers/base/org/jboss/teiid/*

This directory defines JBoss Data Virtualization modules for JBoss EAP.

/modules/system/layers/base/org/jboss/teiid/client

This directory contains JBoss Data Virtualization client libraries. It has the JBoss Data Virtualization JDBC driver JAR file, **teiid-8.4.0.Final-jdbc.jar**, and also contains **teiid-hibernate-dialect-8.4.0.Final.jar** that contains the JBoss Data Virtualization Hibernate dialect.

{standalone, domain}/tmp/teiid

This directory contains temporary files created by JBoss Data Virtualization. These are mostly created by the buffer manager. These files are not needed across a VM restart. Creation of JBoss Data Virtualization LOB values (for example through SQL/XML) will typically create one file per LOB once it exceeds the allowable in memory size of 8KB. In heavy usage scenarios, consider pointing the buffer directory at a partition that is routinely defragmented.

{standalone, domain}/data/teiid-data

This directory contains cached VDB metadata files. Do not edit them manually.

dataVirtualization

This directory contains JDBC drivers, adminshell, the ModeShape VDB and some WARs.

APPENDIX A. REVISION HISTORY

Revision 6.2.0-17
Updates for 6.2.

Thu Dec 10 2015

David Le Sage