

**Teiid - Scalable Information Integration**

**1**

# **Teiid Administrator's Guide**

**7.6**

---

---

---

<b>1. Installation Guide</b> .....	1
1.1. Installation .....	1
1.2. Directory Structure Explained .....	3
1.2.1. bin/scripts .....	3
1.2.2. docs/teiid .....	3
1.2.3. /standalone/configuration/standalone-teiid.xml .....	3
1.2.4. /standalone/deployments .....	3
1.2.5. modules .....	4
1.2.6. /modules/org/jboss/teiid/main/conf .....	4
1.2.7. /domain/configuration/domain-teiid.xml .....	4
1.2.8. modules/org/jboss/teiid/client .....	4
1.2.9. {standlaone or domain}tmp/teiid .....	4
1.2.10. {standlaone or domain}data/teiid-data .....	4
<b>2. Deploying VDBs in Teiid 7</b> .....	5
2.1. Deploying a VDB .....	5
2.1.1. Direct File Deployment .....	5
2.1.2. Admin Console Deployment (Web) .....	5
2.1.3. CLI based Deployment .....	6
2.1.4. AdminShell Deployment .....	6
2.1.5. Admin API Deployment .....	6
2.2. Deploying VDB Dependencies .....	6
2.2.1. JDBC Data Sources .....	7
2.2.2. File Data Sources .....	8
2.2.3. Web Service Data Sources .....	9
2.2.4. Salesforce Data Sources .....	14
2.2.5. LDAP Data Sources .....	15
2.3. VDB Versioning .....	16
2.3.1. Deployment Scenarios .....	17
2.4. Migrating VDBs from 6.x .....	17
<b>3. Teiid Security</b> .....	19
3.1. Authentication .....	19
3.1.1. Pass-through Authentication .....	19
3.2. Authorization .....	19
3.3. Encryption .....	20
3.4. LoginModules .....	20
3.4.1. Built-in LoginModules .....	20
3.4.2. Kerberos support through GSSAPI .....	23
3.4.3. ODBC Client Configuration .....	25
3.4.4. Security at Data Source level .....	25
3.5. Configuring SSL .....	29
3.5.1. SSL Authentication Modes .....	30
3.5.2. Encryption Strength .....	30
<b>4. Logging</b> .....	31
4.1. General Logging .....	31

4.1.1. Logging Contexts .....	31
4.2. Command Logging .....	32
4.3. Audit Logging .....	32
<b>5. Clustering in Teiid .....</b>	<b>35</b>
<b>6. Performance Tuning .....</b>	<b>37</b>
6.1. Memory Management .....	37
6.1.1. Big Data/Memory .....	38
6.1.2. Disk Usage .....	39
6.2. Threading .....	39
6.3. Cache Tuning .....	40
6.4. Socket Transports .....	40
6.5. LOBs .....	40
6.6. Other Considerations .....	41
<b>7. Teiid Admin Console .....</b>	<b>43</b>
7.1. What can be monitored and/or configured? .....	43
7.1.1. Configuration .....	44
7.1.2. Metrics .....	44
7.1.3. Control (Operations) .....	44
7.1.4. Deploying the VDB .....	45
<b>8. AdminShell .....</b>	<b>47</b>
8.1. Introduction .....	47
8.1.1. Download .....	47
8.2. Getting Started .....	47
8.2.1. Essential Rules .....	48
8.2.2. Help .....	49
8.2.3. Basic Commands .....	50
8.3. Executing a script file .....	50
8.4. Log File and Recorded Script file .....	51
8.5. Default Connection Properties .....	51
8.6. Handling Multiple Connections .....	52
8.7. Interactive Shell Nuances .....	53
A. AdminShell Frequently Asked Questions .....	55
B. Other Scripting Environments .....	57
C. System Properties .....	59

# Installation Guide

Starting with the 7.0 release Teiid needs to be installed into an existing JBoss AS installation, which is entirely different from previous versions.



## Note

Teiid does not support the "embedded" mode in 7.6 version. ("embedded" will be coming in a future release).

## 1.1. Installation

Steps to install Teiid

1. Download the [JBoss AS 7.1.0](http://www.jboss.org/jbossas/downloads.html) [http://www.jboss.org/jbossas/downloads.html] application server. Install the server by unzipping into a known location. Ex: /apps/jboss-7.1.0



## Note

You may also choose to use an existing AS installation. However if a previous version of Teiid was already installed, you must remove the old teiid distribution artifacts before installing the new version.

2. Download [Teiid 7.6](http://www.jboss.org/teiid/downloads.html) [http://www.jboss.org/teiid/downloads.html]. Unzip the downloaded artifact inside the JBoss AS installation. Teiid 7.6 directory structure matches JBoss profiles directly - it is just an overlay. This will add necessary modules and configuration files to install Teiid in JBoss AS 7.x. Teiid works in both Standalone and Domain modes. Teiid provides separate configuration files (standalone-teiid.xml and domain-teiid.xml) for both modes apart from the default configuration files come with JBoss AS 7.x

The "Domain" mode recommended in a clustered environment to take advantage of clustered caching and cluster safe distribution of events.

3. Start the JBoss AS server by executing

```
<jboss-install>/bin/standalone.sh --server-config=standalone-teiid.xml
```

if you want to start the "standalone" profile.

To start the server in "Domain" mode, install the JBoss AS 7.x and Teiid on all the servers in are going to be part of the cluster. Select one of the server as the "master" domain controller, the rest of the servers will be slaves that connect to the "master" domain controller for all the administrative operations. Edit "host-teiid.xml" on all the slave servers in "jbossas/domain/configuration" directory, and make sure the "domain-controller" element is configured with ip address of the "master" server. For example

```
<domain-controller>
  <remote host="192.168.100.1" port="9999"/>
</domain-controller>
```

the port number is for the management interface, make sure they match to that of what you have in the configuration. If you want to change the "server-group" name or add unique server names to each server then modify the "servers" element. Once all the configuration edits are done, then you can start JBoss AS in domain mode by executing the command

```
<jboss-install>/bin/domain.sh --domain-config=domain-teiid.xml --host-config=host-teiid.xml
```

on all the servers. Once all the servers are up, to complete the installion for the domain mode, run the following command on any one server. Note that this only needs to be run for once per whole domain install. Since in domain mode you can not statically deploy resources, this script deploys default resources required (file, ldap, salesforce and ws connectors) for Teiid once the server starts using the CLI interface.

```
<jboss-install>/bin/jboss-admin.sh --file=scripts/teiid-domain-mode-install.cli
```

4. That it!. JBoss AS and Teiid are now installed and running. See below instructions to customize various settings.
5. Once VDBs have been deployed, users can now connect their JDBC applications to Teiid. If you need help on connecting your application to the Teiid using JDBC check out the "Client Developer's Guide".

## 1.2. Directory Structure Explained

### Example 1.1. Directory Structure

This shows the contents of the Teiid 7.6 deployment. The directory structure is exactly the same under any JBoss profile.

```
/bin
  /scripts
/docs
  /teiid
    /datsources
    /schema
    /examples
  /domain
  /configuration
    domain-teiid.xml
/modules
  /org/jboss/teiid/*
/standalone
  /configuration
    standalone-teiid.xml
  /deployments
    teiid-*.rar
```

#### 1.2.1. bin/scripts

Has installation and utility CLI scripts.

#### 1.2.2. docs/teiid

Has documents, examples, sample data source XML fragments and schema files. Contains artifacts need by the Quick Start Example.

#### 1.2.3. /standalone/configuration/standalone-teiid.xml

Master configuration file for Teiid system. This file contains Teiid subsystem in addition to standard JBoss AS web profile subsystems

#### 1.2.4. /standalone/deployments

This directory contains all the resource-adapter RAR files that are supplied as part of the Teiid installation. Note that resource adapters are not installed by default for the domain mode.

### 1.2.5. modules

This directory defines the modules for JBoss As system

### 1.2.6. /modules/org/jboss/teiid/main/conf

Relevant Files

- /teiid-security-users.properties
- /teiid-security-roles.properties

These files define the allowed users and their defined roles in Teiid using the default security domain. Edit these files to add uses. If you want to use a different security domain look for details in main configuration file.

### 1.2.7. /domain/configuration/domain-teiid.xml

Master configuration file for Domain mode. This file contains Teiid subsystem in addition to standard JBoss AS web profile subsystems.

### 1.2.8. modules/org/jboss/teiid/client

This directory contains Teiid client libraries. It has the Teiid JDBC driver jar, "teiid-7.6-client.jar", and also contains "teiid-hibernate-dialect-7.6.jar" that contains Teiid's Hibernate dialect.

### 1.2.9. {standlaone or domain}tmp/teiid

This directory contains temporary files created by Teiid. These are mostly created by the buffer manager. These files are not needed across a VM restart. Creation of Teiid 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.

### 1.2.10. {standlaone or domain}data/teiid-data

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

# Deploying VDBs in Teiid 7

A [VDB](http://www.jboss.org/teiid/basics/virtualdatabases.html) [http://www.jboss.org/teiid/basics/virtualdatabases.html] is the primary means to define a Virtual Database in Teiid. A user can create a VDB using [Teiid Designer](http://www.jboss.org/teiid/designer.html) [http://www.jboss.org/teiid/designer.html] or follow the instructions in the Reference Guide to create a "Dynamic VDB" without Teiid Designer.

## 2.1. Deploying a VDB

Once you have a "VDB" built it can be deployed/removed in Teiid runtime in different ways.



### Warning

If *VDB versioning* is not used to give distinct version numbers, overwriting a VDB of the same name will terminate all connections to the old VDB. It is recommended that VDB versioning be used for production systems.



### Note

Removing an existing VDB will immediately clean up VDB file resources, but will not automatically terminate existing sessions.

### 2.1.1. Direct File Deployment

Copy the VDB file into the "`<jboss-install>/standalone/deployments`" directory. Then, create an empty marker file with the same name as the VDB 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". Make sure that there are no other VDB files with the same name. If a VDB already exists with the same name, then this VDB will be replaced with the new VDB. This is the simplest way to deploy a VDB. This is mostly designed for quick deployment during development, when the Teiid server is available locally on the developer machine.



### Note

This only works in the Standalone mode. For domain mode, you can use any other available methods.

### 2.1.2. Admin Console Deployment (Web)

Use the admin web console at:

```
http://<host>:<port>/admin-console
```

More details for this can be found in the [Admin Console VDB deployment section](#). This is the easiest way to deploy a VDB to a remote server.

### 2.1.3. CLI based Deployment

JBoss AS provides command line interface (CLI) for doing any kind of administrative task. Execute "bin/jboss-admin.sh --connect" command and run

```
deploy [--all-server-groups] /path/to/my.vdb
```

to deploy the VDB. Note that in domain mode, you need to either select a particular "server-group" or all available server groups are deployment options. Check out CLI documentation for more general usage of the CLI.

### 2.1.4. AdminShell Deployment

Teiid provides a groovy based AdminShell scripting tool, which can be used to deploy a VDB. Check out the "deploy" method. Consult the [AdminShell documentation](#) for more information. Note that using the AdminShell scripting, you can automate deployment of artifacts in your environment. When using AdminShell, in domain mode, the VDB is deployed to all the servers.

### 2.1.5. Admin API Deployment

The Admin API (look in org.teiid.adminpi.\*) provides Java API methods that lets a user connect to a Teiid runtime and deploy a VDB. If you need to programmatically deploy a VDB use this method. This method is preferable for OEM users, who are trying to extend the Teiid's capabilities through their applications. When using Admin API, in domain mode, the VDB is deployed to all the servers.

## 2.2. Deploying VDB Dependencies

Apart from deploying the VDB, the user is also responsible for providing all the necessary dependent libraries, configuration for creating the data sources that are needed by the models (schemas) defined in "META-INF/vdb.xml" file inside your VDB. For example, if you are trying to integrate data from Oracle RDBMS and File sources in your VDB, then you are responsible for providing the JDBC driver for the Oracle source and any necessary documents and configuration that are needed by the File Translator.

Data source instances may be used by only one VDB, or may be shared with as many VDBs or other applications as makes sense for your deployments. Consider sharing connections to sources that have heavy-weight and resource constrained connections.

With the exception of JDBC sources, other supported data sources have a corresponding JCA connector (.rar) files in "<jboss-install>/standalone/deployments" directory. Either directly edit the standalone-teiid.xml or use CLI or Admin API to create the required data sources by the VDB. Example configurations are provided for all the sources in "<jboss-install>/docs/teiid/datasources" directory. Note that in the domain mode, you must either use CLI or admin-console or AdminShell to configure.

Some data sources may contain passwords or other sensitive information. See the WIKI article [EncryptingDataSourcePasswords](http://community.jboss.org/wiki/EncryptingDataSourcePasswords) [http://community.jboss.org/wiki/EncryptingDataSourcePasswords] to not store passwords in plain text.

Once the VDB and its dependencies are deployed, then client applications can connect using the JDBC API. If there are any errors in the deployment, a connection attempt will not be successful and a message will be logged. You can use the [admin-console](#) tool or check the log files for errors and correct them before proceeding.

### 2.2.1. JDBC Data Sources

The following is an example highlighting configuring an Oracle data source. The process is nearly identical regardless of the database vendor. Typically only the client JDBC jar and the settings like connection url and user credentials change.

There are 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 AS7 is described [here](http://community.jboss.org/docs/DOC-16657) [http://community.jboss.org/docs/DOC-16657].

1. Copy the Oracle JDBC JAR file into "<jboss-install>/standalone/deployments" directory, or use CLI and use "deploy" command to deploy the jar file.
2. Create a "data source" to the Oracle instance in the JBoss container. This typically is done by using "admin-console" application or CLI applications. You can also use Admin API to create this data source. You can manually copy the following xml fragment under "data source" subsystem in standalone-teiid.xml file

```
<datasource jndi-name="java:/OracleDS" pool-name="OracleDS" enabled="true" jta="true"
use-java-context="true" use-ccm="true">
  <connection-url>jdbc:oracle:thin:{host}:1521:orcl</connection-url>
  <driver>ojdbc6.jar</driver>
  <pool>
    <prefill>>false</prefill>
    <use-strict-min>>false</use-strict-min>
    <flush-strategy>FailingConnectionOnly</flush-strategy>
  </pool>
  <security>
```

```
<user-name>{user}</user-name>
<password>{password}</password>
</security>
</datasource>
```

Template files for different databases can be found at `{jboss-as}/docs/teiid/datasources` directory

3. If you like create using CLI, you can issue command like

```
/subsystem=datasources/data-source=oracel-ds:add(jndi-name=java:/OracleDS, pool-
name=oracle-ds, driver-name=oracle, connection-url=jdbc:oracle:thin:{host}:1521:orcl,user-
name={user}, password={password})
/subsystem=datasources/data-source=oracel-ds:enable
```

### 2.2.2. File Data Sources

File data sources use a Teiid specific JCA connector. You need to create following XML fragment, and copy it under "resource-adapters" subsystem in standalone.xml file or use the CLI to create connection factory.

#### Example 2.1. Template for creating a File based data source

```
<!-- If subsystem is already defined, only copy the contents under it and edit to suit your needs -->
<subsystem xmlns="urn:jboss:domain:resource-adapters:1.0">
  <resource-adapters>
    <resource-adapter>
      <archive>teiid-connector-file.rar</archive>
      <transaction-support>NoTransaction</transaction-support>
      <connection-definitions>
        <connection-definition class-
name="org.teiid.resource.adapter.file.FileManagedConnectionFactory"
          jndi-name="java:/fileDS"
          enabled="true"
          use-java-context="true"
          pool-name="file-ds">
          <!-- Directory where the data files are stored -->
          <config-property name="ParentDirectory"/>/home/rareddy/testing/</config-property>
        <!-- Optional properties -->
```

```

        <!-- 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)*
            -->
        <!-- <config-property name="FileMapping">file1.txt=fileX.txt,file2.txt=fileY.txt</config-
        property> -->

        <!-- Set AllowParentPaths to false to disallow .. in paths.
            This prevent requesting files that are not contained in the parent directory -->
        <config-property name="AllowParentPaths">true</config-property>
    </connection-definition>
</connection-definitions>
</resource-adapter>
</resource-adapters>
</subsystem>

```

You can use CLI to configure the same, however it is a multi-step process. For domain mode, you must use CLI, or admin-console or AdminShell/Admin API

### 2.2.3. Web Service Data Sources

Web service data sources use a Teiid specific JCA connector. You need to create following XML fragment, and copy it under "resource-adapters" subsystem in standalone.xml file or use the CLI to create connection factory.

#### Example 2.2. Template for creating a web service based data source

```

<!-- If subsystem is already defined, only copy the contents under it and edit to suit your needs -->
<subsystem xmlns="urn:jboss:domain:resource-adapters:1.0">
  <resource-adapters>
    <resource-adapter>
      <archive>teiid-connector-ws.rar</archive>
      <transaction-support>NoTransaction</transaction-support>
      <connection-definitions>
        <connection-definition class-
name="org.teiid.resource.adapter.ws.WSManagedConnectionFactory"
          jndi-name="java:/wsDS"
          enabled="true"
          use-java-context="true"
          pool-name="ws-ds">
          <config-property name="EndPoint">http://somewhere.com</config-property>
        </connection-definition>
      </connection-definitions>
    </resource-adapter>
  </resource-adapters>
</subsystem>

```

```
</resource-adapter>
</resource-adapters>
</subsystem>
```

You can use CLI to configure the same, however it is a multi-step process. For domain mode, you must use CLI, or admin-console or AdminShell/Admin API

### 2.2.3.1. CXF Configuration

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 1 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 be `http://teiid.org`. See the sections on WS-Security, Logging, etc. for examples of using the CXF configuration file.

### Example 2.3. Sample Spring XML Configuration To Set 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/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}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`, is not configurable. The local name is followed by `.http-conduit`. It will be based upon the `configName` setting, with a default value of `teiid`.

See the [CXF documentation](http://cxf.apache.org/docs/) [http://cxf.apache.org/docs/] for all possible configuration options.



#### Note

The CXF configuration is currently only applicable to non-binary web service calls.

### 2.2.3.2. WS-Security

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.

#### Example 2.4. Example WS-Security enabled data source

```
<!-- If subsystem is already defined, only copy the contents under it and edit to suit your needs -->
<subsystem xmlns="urn:jboss:domain:resource-adapters:1.0">
  <resource-adapters>
    <resource-adapter>
      <archive>teiid-connector-ws.rar</archive>
      <transaction-support>NoTransaction</transaction-support>
      <connection-definitions>
        <connection-definition class-
name="org.teiid.resource.adapter.ws.WSManagedConnectionFactory"
          jndi-name="java:/wsDS"
          enabled="true"
          use-java-context="true"
          pool-name="ws-ds">
          <config-property name="EndPoint">http://somewhere.com</config-property>
          <config-property name="ConfigFile">${jboss.server.home.dir}/server/default/conf/
xxx-jbossws-cxf.xml</config-property>
          <config-property name="ConfigName">port_x</config-property>
          <config-property name="SecurityType">WSSecurity</config-property>
        </connection-definition>
      </connection-definitions>
    </resource-adapter>
  </resource-adapters>
</subsystem>
```

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 class="org.apache.cxf.binding.soap.saaj.SAAJOutInterceptor"/>
        <ref bean="Timestamp_Request"/>
      </jaxws:outInterceptors>
    </jaxws:client>

    <bean
      class="org.apache.cxf.ws.security.wss4j.WSS4JOutInterceptor"
      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. The configuration may contain other port configurations with different local names.

For more information on configuring CXF interceptors, please consult the [CXF documentation](https://cwiki.apache.org/CXF20DOC/ws-security.html) [https://cwiki.apache.org/CXF20DOC/ws-security.html] or the [JBossWS-CXF documentation](http://community.jboss.org/wiki/JBossWS-StackCXFUserGuide#WSSecurity) [http://community.jboss.org/wiki/JBossWS-StackCXFUserGuide#WSSecurity].

### 2.2.3.3. Logging

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.

### Example 2.5. Example logging data source

```
<!-- If subsystem is already defined, only copy the contents under it and edit to suit your needs -->
<subsystem xmlns="urn:jboss:domain:resource-adapters:1.0">
  <resource-adapters>
    <resource-adapter>
      <archive>teiid-connector-ws.rar</archive>
      <transaction-support>NoTransaction</transaction-support>
```

```

<connection-definitions>
    <connection-definition class-
name="org.teiid.resource.adapter.ws.WSManagedConnectionFactory"
    jndi-name="java:/wsDS"
    enabled="true"
    use-java-context="true"
    pool-name="ws-ds">
        <config-property name="EndPoint">http://somewhere.com</config-property>
        <config-property name="ConfigFile">${jboss.server.home.dir}/server/default/conf/
xxx-jbossws-cxf.xml</config-property>
        <config-property name="ConfigName">port_y</config-property>
    </connection-definition>
</connection-definitions>
</resource-adapter>
</resource-adapters>
</subsystem>

```

Corresponding xxx-jbossws-cxf.xml

```

<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>

```

#### 2.2.3.4. Transport Settings

The CXF config property may also be used to control low level aspects of the HTTP transport. See the [CXF documentation](http://cxf.apache.org/docs/client-http-transport-including-ssl-support.html) [http://cxf.apache.org/docs/client-http-transport-including-ssl-support.html] for all possible options.

**Example 2.6. Example Disabling Hostname Verification**

```

<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>

```

**2.2.4. Salesforce Data Sources**

Salesforce data sources use a Teiid specific JCA connector. You need to create following XML fragment, and copy it under "resource-adapters" subsystem in standalone.xml file or use the CLI to create connection factory.

**Example 2.7. Template for creating a Salesforce based data source**

```

<!-- If subsystem is already defined, only copy the contents under it and edit to suit your needs -->
<subsystem xmlns="urn:jboss:domain:resource-adapters:1.0">
  <resource-adapters>
    <resource-adapter>
      <archive>teiid-connector-salesforce.rar</archive>
      <transaction-support>NoTransaction</transaction-support>
      <connection-definitions>
        <connection-definition class-
name="org.teiid.resource.adapter.salesforce.SalesForceManagedConnectionFactory"
          jndi-name="java:sfDS"
          enabled="true"
          use-java-context="true"
          pool-name="sf-ds">
          <config-property name="URL">https://test.salesforce.com/services/Soap/u/10.0</
config-property>
          <config-property name="username">username</config-property>

```

```

        <config-property name="password">password</config-property>
    </connection-definition>
</connection-definitions>
</resource-adapter>
</resource-adapters>
</subsystem>

```

You can use CLI to configure the same, however it is a multi-step process. For domain mode, you must use CLI, or admin-console or AdminShell/Admin API

### 2.2.4.1. CXF Configuration

Salesforce 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 1 port configuration can be used by this data source. The namespace URI for the QName in your config file should be "urn:partner.soap.sforce.com", with configuration name "SforceService". For sample cxf configuration file and details on configuration see [Web Service Data Sources](#)

See the [CXF documentation](http://cxf.apache.org/docs/) [http://cxf.apache.org/docs/] for all possible configuration options.



#### Note

The CXF configuration in Salesforce data source is only used for http bus configuration not for purposes of ws-security, Salesforce has its own security authentication.

### 2.2.5. LDAP Data Sources

LDAP data sources use a Teiid specific JCA connector. You need to create following XML fragment, and copy it under "resource-adapters" subsystem in standalone.xml file or use the CLI to create connection factory.

#### Example 2.8. Template for creating an LDAP based data source

```

<!-- If subsystem is already defined, only copy the contents under it and edit to suit your needs -->
<subsystem xmlns="urn:jboss:domain:resource-adapters:1.0">
  <resource-adapters>
    <resource-adapter>
      <archive>teiid-connector-ldap.rar</archive>
      <transaction-support>NoTransaction</transaction-support>
      <connection-definitions>

```

```

                                                                 <connection-definition   class-
name="org.teiid.resource.adapter.Ldap.LDAPManagedConnectionFactory"
    jndi-name="java:/ldapDS"
    enabled="true"
    use-java-context="true"
    pool-name="ldap-ds">

    <config-property name="LdapAdminUserDN">cn=x,ou=y,dc=z</config-property>
    <config-property name="LdapAdminUserPassword">password</config-property>
    <config-property name="LdapUrl">ldap://ldapServer:389</config-property>
  </connection-definition>
</connection-definitions>
</resource-adapter>
</resource-adapters>
</subsystem>
```

You can use CLI to configure the same, however it is a multi-step process. For domain mode, you must use CLI, or admin-console or AdminShell/Admin API

### 2.3. VDB Versioning

VDB Versioning is a feature that allows multiple versions of a VDB to be deployed at the same time with additional support to determine which version will be used. When a user connects to Teiid the desired VDB version can be set as a connection property (See the Client Developers Guide). If a specific version is set, then only that VDB may be connected to. If no version is set, then the deployed VDBs are searched for the appropriate version. This feature helps support more fluid migration scenarios.

Setting the version can either be done in the vdb.xml, which is useful for dynamic vdb's, or through a naming convention of the deployment file - vdbname.version.vdb, e.g. marketdata.2.vdb. The deployer is responsible for choosing an appropriate version number. If there is already a VDB name/version 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.

Once deployed a VDB has an updatable property called connection type, which is used 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 vdb's marked as *ANY*.
- *ANY* - allow connections with or without a version specified.

The connection type may be changed either through the AdminConsole or the AdminAPI.

### 2.3.1. Deployment Scenarios

If only a select few applications are to migrate to the new VDB version, then a freshly deployed VDB would be left as BY\_VERSION. This ensures that only applications that know the new version may use it.

If only a select few applications are to remain on the current VDB version, then their connection settings would need to be updated to reference the current VDB by its version. Then the newly deployed vdb would have its connection type set to ANY, which allows all new connections to be made against the newer version. If a rollback is needed in this scenario, then the newly deployed vdb would have its connection type set to NONE or BY\_VERSION accordingly.

## 2.4. Migrating VDBs from 6.x

VDBs from prior release contain an older configuration file version that is no longer supported. Note - XML and File based sources from previous releases have changed, and require manual changes to the VDB.



# Teiid Security

The Teiid system provides a range of built-in and extensible security features to enable the secure access of data.

## 3.1. Authentication

JDBC clients may use simple passwords to authenticate a user.

Typically a user name is required, however user names may be considered optional if the identity of the user can be discerned by the password credential alone. In any case it is up to the configured security domain to determine whether a user can be authenticated. If you need authentication, the administrator must configure a LoginModule to be used with Teiid. See below for more information on how configure the Login module in JBoss AS.



### Note

By default, access to Teiid is NOT secure. The default login modules are only backed by file based authentication, which has a well known user name and password. The same is true for making connections to the Admin Console application. We DO NOT recommend leaving the default security profile as defined when you are exposing sensitive data.

### 3.1.1. Pass-through Authentication

If your client application (web application or Web service) resides in the same JBoss AS instance as Teiid and the client application uses a security-domain, then you can configure Teiid to use the same security-domain and not force the user to re-authenticate. In this case Teiid looks for an authenticated subject in the calling thread context and uses it for sessioning and authorization. To configure Teiid for pass-through authentication, change the Teiid security-domain name to the same name as your application's security domain name in the "standalone-teiid.xml" file in the "transport" section. Please note that for this to work, the security-domain must be a JAAS based Login Module and your client application MUST obtain its Teiid connection using a *local/embedded* connection with the *PassthroughAuthentication=true* connection flag set.

## 3.2. Authorization

Authorization covers both administrative activities and data roles. A data role is a collection of permissions (also referred to as entitlements) and a collection of entitled principals or groups. With the deployment of a VDB the deployer can choose which principals and groups have which data roles.

### 3.3. Encryption

At a transport level Teiid provides built-in support for JDBC over SSL or just sensitive message encryption when SSL is not in use.

Passwords in configuration files however are by default stored in plain text. If you need these values to be encrypted, please see [encrypting passwords](http://community.jboss.org/wiki/maskingpasswordsinjbossasxmlconfiguration) [http://community.jboss.org/wiki/maskingpasswordsinjbossasxmlconfiguration] for instructions on encryption facilities provided by the container.

### 3.4. LoginModules

LoginModules are an essential part of the JAAS security framework and provide Teiid customizable user authentication and the ability to reuse existing LoginModules defined for JBossAS. See [JBossAS Security](http://docs.jboss.org/jbossas/admindevel326/html/ch8.chapter.html) [http://docs.jboss.org/jbossas/admindevel326/html/ch8.chapter.html] for general information on configuring security in JBossAS.

Teiid can be configured with multiple named application policies that group together relevant LoginModules. Each of these application policy (or domain) names can be used to fully qualify user names to authenticate only against that domain. The format for a qualified name is `username@domainname`. The `security-domain` attribute under the `transport/authentication` element in the `<jboss-install>/standalone/configuration/standalone-teiid.xml` file should be used set the comma separated list of desired domains. The `securityDomains` property defaults to a single domain named `teiid-security`.

If a user name is not fully qualified, then the installed domains will be consulted in order until a domain successfully or unsuccessfully authenticates the user.

If no domain can authenticate the user, the login attempt will fail. Details of the failed attempt including invalid users, which domains were consulted, etc. will be in the server log with appropriate levels of severity.



#### Note

The `security-domain` defined for each transport type is can be different. The JDBC connection . The default name of JDBC connection's `security-domain` is "teiid-security". In existing installations an appropriate security domain may already be configured for use by administrative clients (typically "admin-console"). If the admin connections (CLI and adminshell) are not secured, it is recommended that you secure that that interface by executing one of scripts in the "bin/scripts" directory.

#### 3.4.1. Built-in LoginModules

JBossAS provides several LoginModules for common authentication needs, such as authenticating from a [Section 3.4.1.1, "Text Based LoginModule"](#) or a [Section 3.4.1.2, "LDAP Based LoginModule"](#).

You can install multiple login modules as part of single security domain configuration and configure them to part of login process. For example, for "teiid-security" domain, you can configure a file based and also LDAP based login modules, and have your user authenticated with either or both login modules. If you want to write your own custom login module, refer to the Developer's Guide for instructions.

### 3.4.1.1. Text Based LoginModule

The UsersRolesLoginModule utilizes simple text files to authenticate users and to define their groups. The below XML fragment under "security" subsystem shows a Text based login module. Also see standalone-teiid.xml configuration file contains an example of how to use UsersRolesLoginModule.

```
<subsystem xmlns="urn:jboss:domain:security:1.1">
  <security-domains>
    <security-domain name="teiid-security" cache-type="default">
      <authentication>
        <login-module code="UsersRoles" flag="required">
          <module-option name="usersProperties" value="teiid-security-users.properties"/>
          <module-option name="rolesProperties" value="teiid-security-roles.properties"/>
        </login-module>
      </authentication>
    </security-domain>
  </security-domains>
</subsystem>
```



#### Note

The UsersRolesLoginModule is not recommended for production use and is strongly recommended that you replace this login module.

User names and passwords are stored in the <jboss-as>/modules/org/jboss/teiid/main/conf/teiid-security-users.properties file. These files must be available on classpath

#### Example 3.1. Example user.properties file

```
# A users.properties file for use with the UsersRolesLoginModule
# username=password

fred=password
george=password
```

...

JAAS role assignments are stored in the `<jboss-as>/modules/org/jboss/teiid/main/conf/teiid-security-roles.properties` file.

### Example 3.2. Example user.properties file

```
# A roles.properties file for use with the UsersRolesLoginModule
# username=role1,role2,...

data_role_1=fred,sally
data_role_2=george
```

User and role names are entirely up to the needs of the given deployment. For example each application team can set their own security constraints for their VDBs, by mapping their VDB data roles to application specific JAAS roles, e.g. `app_role_1=user1,user2,user3`.



#### Note

Teiid data roles names are independent of JAAS roles. VDB creators can choose whatever name they want for their data roles, which are then mapped at deployment time to JAAS roles.

### 3.4.1.2. LDAP Based LoginModule

See [LDAP LoginModule configuration](http://community.jboss.org/docs/DOC-11253) [http://community.jboss.org/docs/DOC-11253] for the AS community guide. The following are streamlined installation instruction.

1. If using SSL to the LDAP server, ensure that the Corporate CA Certificate is added to the JRE trust store.
2. Include LDAP LoginModule in the JAAS Configuration

Configure LDAP authentication by editing `standalone-teiid.xml` under security subdomain. If you wish to configure specifically for teiid, then the security domain `teiid-security` will need to be created/alterd. You could do one of the following for Teiid:

- Reuse the `admin-console` (or whatever name you choose) security domain for Teiid by changing the teiid configuration `<jboss-install>/standalone/configuration/standalone-teiid.xml` to point to `admin-console`.
- Follow the same steps to configure an LDAP security domain named `teiid-security`.

3. Obscure the LDAP Password

Finally, protect the password following [these instructions](http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html/Security_Guide/Using_LdapExtLoginModule_with_JaasSecurityDomain.html). [http://docs.redhat.com/docs/en-US/JBoss\_Enterprise\_Application\_Platform/5/html/Security\_Guide/Using\_LdapExtLoginModule\_with\_JaasSecurityDomain.html] Note that the salt must be 8 chars and see also <http://community.jboss.org/message/137756#137756> for more on securing passwords.

### 3.4.2. Kerberos support through GSSAPI

Teiid supports kerberos authentication using GSSAPI, to be used with single sign-on applications. This service ticket negotiation based authentication is supported through remote JDBC and ODBC drivers and LocalConnections. Client configuration is different for based on connection you are using

#### 3.4.2.1. LocalConnection

Set the JDBC URL property *PassthroughAuthentication* as true and use [JBoss Negotiation](http://community.jboss.org/docs/DOC-10680) [http://community.jboss.org/docs/DOC-10680] for authentication of your web-application with kerberos. When the web application authenticates with the provided kerberos token, the same subject authenticated will be used in Teiid. For details about configuration, check the JBoss Negotiation documentation.

#### 3.4.2.2. Remote Connections

On the server, edit the `<jboss-install>/standalone/configuration/standalone-teiid.xml` under teiid subsystem on "transport" definition, add follows:

```
<transport name="jdbc" protocol="teiid" socket-binding="teiid-jdbc"/>
  <authentication security-domain="teiid-security" krb5-domain="krb5-domain"/>
</transport>
```

Now we need to define a security domain context for kerberos with the name mentioned (krb5-domain)in above. Since kerberos authorization cannot define authorization roles, we'll define them using another login context. Given below is a sample configuration to define roles using a UserRolesLoginModule.



#### Note

This configuration replaces the default Teiid login configuration, and you should change the principal and key tab locations accordingly.

```
<!--login module that negotiates the login context for kerberos -->
```

```
<subsystem xmlns="urn:jboss:domain:security:1.1">
  <security-domains>
    <security-domain name="krb5-domain" cache-type="default">
      <authentication>
        <login-module code="Kerberos" flag="required">
          <module-option name="storeKey">true</module-option>
          <module-option name="useKeyTab">true</module-option>
          <module-option name="principal">demo@EXAMPLE.COM</module-option>
          <module-option name="keyTab">path/to/krb5.keytab</module-option>
          <module-option name="doNotPrompt">true</module-option>
          <module-option name="debug">false</module-option>
        </login-module>
      </authentication>
    </security-domain>
    <!-- teiid's default security domain, replace this with your own if needs to be any other
JAAS domain -->
    <security-domain name="teiid-security" cache-type="default">
      <authentication>
        <login-module code="UsersRoles" flag="required">
          <module-option name="usersProperties" value="teiid-security-users.properties"/>
          <module-option name="rolesProperties" value="teiid-security-roles.properties"/>
        </login-module>
      </authentication>
    </security-domain>
  </security-domains>
</subsystem>
```

Edit the "standalone.conf" file in the "\${jboss-as}/bin" directory and add the following JVM options (changing the realm and KDC settings according to your environment)

```
JAVA_OPTS = "$JAVA_OPTS -Djava.security.krb5.realm=EXAMPLE.COM -
Djava.security.krb5.kdc=kerberos.example.com -
Djavax.security.auth.useSubjectCredsOnly=false"
```

This finishes the configuration on the server side, restart the server and make sure that there were no errors during startup.

### 3.4.2.2.1. JDBC Client Configuration

In you client VM the JAAS configuration for kerberos authentication needs to be written. A sample configuration file (client.conf) is show below

```
Client {
  com.sun.security.auth.module.Krb5LoginModule required
  useTicketCache=true
  storeKey=true
  useKeyTab=true
  keyTab="/path/to/krb5.keytab"
  doNotPrompt=false
  debug=false
  principal="demo@EXAMPLE.COM";
};
```

Add the following JVM options to your client's startup script - change Realm and KDC settings according to your environment

```
-Djava.security.krb5.realm=EXAMPLE.COM
-Djava.security.krb5.kdc=kerberos.example.com
-Djavax.security.auth.useSubjectCredsOnly=false
-Dsun.security.krb5.debug=false
-Djava.security.auth.login.config=/path/to/client.conf
```

Add the following URL connection properties to Teiid JDBC connection string

```
authenticationType=KRB5;jaasName=Client;kerberosServicePrincipleName=demo@EXAMPLE.COM
```

There is no need to provide the user name and password. When the application makes a JDBC connection, it will authenticate locally and use the same user credentials to negotiate a service token with server and grant the connection. See Client Developer's guide for information on connection properties and how to configure data sources.

### 3.4.3. ODBC Client Configuration

Consult the PostgreSQL ODBC client documentation.

### 3.4.4. Security at Data Source level

In some use cases, the user might need to pass-in different credentials to their data sources based on the logged in user rather than using the shared credentials for all the logged users. To support this feature, JBoss AS and Teiid provide multiple different login modules to be used in conjunction with Teiid's main security domain. See this [document](http://community.jboss.org/docs/DOC-9350) [http://community.jboss.org/docs/DOC-9350] for details on configuration. Note that the below directions need to be used in conjunction with this document.

### 3.4.4.1. CallerIdentity and Trusted Payload

If client wants to pass in simple text password or a certificate or a custom serialized object as token credential to the data source, user can configure "CallerIdentity" login module. Using this login module, user can pass-in same credential that user logged into Teiid security domain to the data source. Here is a sample configuration, this needs to be configured in "standalone-teiid.xml" file.

```
<subsystem xmlns="urn:jboss:domain:security:1.1">
  <security-domains>
    <security-domain name="my-security-domain" cache-type="default">
      <authentication>
        <login-module code="UsersRoles" flag="required">
          <module-option name = "password-stacking">useFirstPass</module-option>
          <module-option name="usersProperties">props/teiid-security-users.properties</
module-option>
          <module-option name="rolesProperties">props/teiid-security-roles.properties</
module-option>
        </login-module>

        <login-module code="org.jboss.resource.security.CallerIdentityLoginModule"
flag="required">
          <module-option name = "password-stacking">useFirstPass</module-option>
          <module-option
name = "managedConnectionFactoryName">jboss.jca:service=LocalTxCM,name=DefaultDS</
module-option>
        </login-module>

      </authentication>
    </security-domain>
  </security-domains>
```

In the datasource defined as the "managedConnectionFactoryName" in the above configuration, you need to add the following element

```
<security-domain>teiid-security</security-domain>
```

In the above configuration example, in the primary login module "UsersRoles" is setup to hold the passwords in the file, and when user logs in with password, the same password will be also set on the logged in Subject after authentication. These credentials can be extracted by the data source by asking for Subject's private credentials.

To use a certificate or serialized object instead of plain password as the token, replace the simple text password with Base64 encoded contents of the serialized object. Please note that encoding and decoding of this object is strictly up to the user as JBoss AS and Teiid will only act as a carrier of the information from login module to connection factory. Using this CallerIdentity module, the connection pool for data source is segmented by Subject.

### 3.4.4.2. Role Based Credential Map

In some use cases, the users are divided by their functionality and they have varied levels of security access to data sources. These types of users are identified by their roles as to what they have access to. In the above "CallerIdentity" login scenario, that may be too fine-grained security at data sources, that can lead resource exhaustion as every user has their own separate connection. Using Role based security gives a balance, where the users with same role are treated equally for authentication purposes at the data source. Teiid provides a login module called "RoleBasedCredentialMap" for this purposes, where administrator can define a role-based authentication module, where given the role of the user from the primary login module, this module will hold a credential to that role. So, it is the container of credentials that maps to different roles. If a user has multiple roles, the first role that has the credential will be chosen. Below find the sample configuration.

```
<subsystem xmlns="urn:jboss:domain:security:1.1">
  <security-domains>
    <security-domain name="my-security-domain" cache-type="default">
      <authentication>

        <login-module code="UsersRoles" flag="required">
          <module-option name = "password-stacking">useFirstPass</module-option>
          <module-option name="usersProperties">teiid-security-users.properties</
module-option>
          <module-option name="rolesProperties">teiid-security-roles.properties</module-
option>
        </login-module>

        <login-module code="org.teiid.jboss.RoleBasedCredentialMapIdentityLoginModule"
flag="required">
          <module-option name = "password-stacking">useFirstPass</module-option>
          <module-option name="credentialMap">teiid-credentialmap.properties</module-
option>
          <module-option
name = "managedConnectionFactoryName">jboss.jca:service=LocalTxCM,name=DefaultDS</
module-option>
        </login-module>

      </authentication>
    </security-domain>
  </security-domains>
</subsystem>
```

```
</security-domain>
</security-domains>
```

In the datasource that is defined as the "managedConnectionFactoryName" in the above configuration, you need to add the following element

```
<security-domain>teiid-security</security-domain>
```

In the above configuration example, in the primary login module "UsersRolesLoginModule" is set up for logging in the primary user and assign some roles. The "RoleBasedCredentialMap" login module is configured to hold role to password information in the file defined by "credentialMap" property. When user logs in, the role information from the primary login module is taken, and extracts the role's password and attaches as a private credential to the Subject. If you want use this for role based trusted token, you can configure the Base64 based encoding/decoded object as defined above.

You can also encrypt the password instead of plain text password using this module. Just include the encrypted password in the file defined by the "credentialMap" property, and define following properties in the "RoleBasedCredentialMap" login module.

```
<login-module          code="org.teiid.jboss.RoleBasedCredentialMapIdentityLoginModule"
flag="required">
  <module-option name = "password-stacking">useFirstPass</module-option>
  <module-option name="credentialMap">props/teiid-credentialmap.properties</module-option>
                                <module-option          name          =
option>

  <!-- below properties are only required when passwords are encrypted -->
  <module-option name = "pbealgo">PBESWithMD5AndDES</module-option>
  <module-option name = "pbepass">testPBESWithMD5AndDES</module-option>
  <module-option name = "salt">abcdefgh</module-option>
  <module-option name = "iterationCount">19</module-option>
</login-module>
```

For full details about encryption of the password, please follow this [document](http://community.jboss.org/docs/DOC-9703) [http://community.jboss.org/docs/DOC-9703]'s "A KeyStore based login module for encrypting a datasource password" section. Be sure to give the same configuration elements in the above configuration, as they are used to encrypt the password.

## 3.5. Configuring SSL

The Teiid's configuration file(s) `standalone-teiid.xml/domain-teiid.xml`, contain transports defined for access to Teiid, and transport contain the properties to configure SSL for socket per transport.

There are two types of transports, each with it's own SSL configuration:

- "teiid" - Default configuration to only encrypt login traffic, none of the other properties are used.
- "pg" type - Defaults to no SSL.

### Example 3.3. Example Configuration

```
<ssl mode="login" authentication-mode="1-way" ssl-protocol="SSLv3" keymanagement-
algorithm="algo"
                                enabled-cipher-
suites="SSL_RSA_WITH_RC4_128_MD5,SSL_RSA_WITH_RC4_128_SHA">
  <keystore name="cert.keystore" password="passwd" type="JKS"/>
  <truststore name="cert.truststore" password="passwd"/>
</ssl>
```

### Properties

- mode - disabled|login|enabled, disabled = no transport or message level security will be used. login = only the login traffic will be encrypted at a message level using 128 bit AES with an ephemeral DH key exchange. No other config values are needed in this mode. enabled = traffic will be secured using the other configuration properties.
- ssl-protocol- Type of SSL protocol to be used. Default is TLSv1
- keystore/type - Keystore type created by the keytool. Default "JKS" is used.
- authentication-mode - anonymous|1-way|2-way, Type of [SSL Authentication Mode](#).
- keymanagement-algorithm - Type of key algorithm used. Default is based upon the VM, e.g. "SunX509"
- keystore/name - The file name of the keystore, which contains the private key of the Server. The file name can be relative resource path available to the Teiid deployer classloader or an absolute file system path. A typical installation would place the keystore file in the conf directory of the profile where Teiid is deployed with a file name relative to the conf path.
- keystorePassword - password for the keystore.

- `truststore/name` - if "authenticationMode" is chosen as "2-way", then this property must be provided. This is the truststore that contains the public key for the client. Depending upon how you created the keystore and truststores, this may be same file as defined under "keystoreFilename" property.
- `truststore/password` - password for the truststore.
- `enabled-cipher-suites` - A comma separated list of cipher suites allowed for encryption between server and client. The values must be valid supported cipher suites otherwise SSL connections will fail.

### 3.5.1. SSL Authentication Modes

SSL supports multiple authentication modes. In most secure intranet environments, anonymous is suitable to just bulk encrypt traffic without the need to setup SSL certificates.

- *anonymous* - no certificates are exchanged, settings are not needed for the keystore and truststore properties. Client must have `org.teiid.ssl.allowAnon` set to true (the default) to connect to an anonymous server.
- *1-way* - the server will present a certificate, which is obtained from the keystore related properties. The client should have a truststore configured to accept the server certificate.
- *2-way* - the server will present a certificate, which is obtained from the keystore related properties. The client should have a truststore configured to accept the server certificate. The client is also expected to present a certificate, which is obtained from its keystore. The client certificate should be accepted by the trust store configured by the truststore related properties.

### 3.5.2. Encryption Strength

Both anonymous SSL and login only encryption are configured to use 128 bit AES encryption by default. By default, 1-way and 2-way SSL allow for cipher suite negotiation based upon the default cipher suites supported by the respective Java platforms of the client and server. User can restrict the cipher suites used for encryption by specifying the `enabledCipherSuites` property above in ssl configuration.

# Logging

## 4.1. General Logging

The Teiid system provides a wealth of information via logging. To control logging level, contexts, and log locations, you should be familiar with [log4j](http://logging.apache.org/log4j/) [http://logging.apache.org/log4j/] and the container's standalone-teiid/domin-teiid.xml configuration file.

All the logs produced by Teiid are prefixed by "org.teiid". This makes it extremely easy to control of of Teiid logging from a single context. Note however that changes to the log configuration file require a restart to take affect

### 4.1.1. Logging Contexts

While all of Teiid's logs are prefixed with "org.teiid", there are more specific contexts depending on the functional area of the system. Note that logs originating from third-party code, including integrated org.jboss components, will be logged through their respective contexts and not through org.teiid. See the table below for information on contexts relevant to Teiid. See the container's standalone-teiid.xml for a more complete listing of logging contexts used in the container.

Context	Description
com.arjuna	Third-party transaction manager. This will include information about all transactions, not just those for Teiid.
org.teiid	Root context for all Teiid logs. Note: there are potentially other contexts used under org.teiid than are shown in this table.
org.teiid.PROCESSOR	Query processing logs. See also org.teiid.PLANNER for query planning logs.
org.teiid.PLANNER	Query planning logs.
org.teiid.SECURITY	Session/Authentication events - see also AUDIT logging
org.teiid.TRANSPORT	Events related to the socket transport.
org.teiid.RUNTIME	Events related to work management and system start/stop.
org.teiid.CONNECTOR	Connector logs.
org.teiid.BUFFER_MGR	Buffer and storage management logs.
org.teiid.TXN_LOG	Detail log of all transaction operations.
org.teiid.COMMAND_LOG	See <a href="#">command logging</a>
org.teiid.AUDIT_LOG	See <a href="#">audit logging</a>

Context	Description
org.teiid.ADMIN_API	Admin API logs.
org.teiid.ODBC	ODBC logs.

### 4.2. Command Logging

Command logging captures executing commands in the Teiid System. Both user commands (that have been submitted to Teiid) and data source commands (that are being executed by the connectors) are tracked through command logging.

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="INFO"/>
  <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="INFO"/>
  <handlers>
    <handler name="COMMAND_FILE"/>
  </handlers>
</logger>
```

See the Developer's Guide to develop a custom logging solution if file based, or any other built-in Log4j, logging is not sufficient.

### 4.3. Audit Logging

Audit logging captures important security events. This includes the enforcement of permissions, authentication success/failures, etc.

To enable audit logging to the default log location, simply enable the `DETAIL` level of logging for the `org.teiid.AUDIT_LOG` context.

To enable audit logging to an alternative file location, configure a separate file appender for the `DETAIL` logging of the `org.teiid.AUDIT_LOG` context. See the Developer's Guide to develop a custom logging solution if file based, or any other built-in Log4j, logging is not sufficient.



# Clustering in Teiid

Since Teiid is installed in JBoss AS, there is no separate configuration needed on the part of the user to cluster the Teiid instances. Check the [installation](#) for quick instructions. See detailed JBoss AS [instructions](https://docs.jboss.org/author/display/AS7/Documentation) [https://docs.jboss.org/author/display/AS7/Documentation] here. Just make sure that you installed Teiid in every JBoss AS node and started JBoss AS instances in the Domain mode before starting the cluster.

Typically users create clusters to improve the performance of the system through:

1. Load Balancing: Take look at the Client developers guide on how to use load balancing between multiple nodes.
2. Fail Over: Take look at the Client developers guide on how to use fail over between multiple nodes.
3. Distributed Caching: This is automatically done for you once you configure it as specified above.
4. Event distribution: metadata and data modifications will be distributed to all cluster members.

In the domain mode, the only way user can deploy any artifacts is using either CLI or using the Admin API or Admin Shell. Copying VDB directly into the "deployments" directory is not supported.



# Performance Tuning

## 6.1. Memory Management

The BufferManager is responsible for tracking both memory and disk usage by Teiid. Configuring the BufferManager properly is one of the most important parts of ensuring high performance. See the file for all BufferManager settings.

The Teiid engine uses batching to reduce the number of memory rows processed at a given time. The batch sizes may be adjusted to larger values if few clients will be accessing the Teiid server simultaneously.

The `max-reserve-kb` setting determines the total size in kilobytes of batches that can be held by the BufferManager 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 auto-calculate a typical max based upon the max heap available to the VM. The auto-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 the AS and other Teiid purposes) and 75% of the memory beyond that.

The BufferManager automatically triggers the use of a canonical value cache if enabled when more than 25% of the reserve is in use. This can dramatically cut the memory usage in situations where similar value sets are being read through Teiid, but does introduce a lookup cost. If you are processing small or highly similar datasets through Teiid, 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 upon actual column values and whether value caching is enabled. Large non built-in type objects can exceed their default size estimate. If an out of memory errors occur, then set a lower the `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.

The `max-processing-kb` setting determines the total size in kilobytes of batches that can be used by active plans regardless of the memory held based on `max-reserve-kb`. The default value of -1 will auto-calculate a typical max based upon the max heap available to the VM and max active plans. The auto-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 the AS and other Teiid purposes).

In systems where large intermediate results are normal (scrolling cursors or sorting over millions of rows) you can consider increasing the `max-processing-kb` and decreasing the `max-reserve-kb` so that each request has access to an effectively smaller buffer space.

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 `max-file-size` megabytes. Consider increasing the storage space available to all such files `max-buffer-space` if your installation makes use of internal materialization, makes heavy use of SQL/XML, or processes large row counts.

### 6.1.1. Big Data/Memory

Usage of extremely large VM sizes and or datasets requires additional considerations. Teiid 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 dealing with datasets with billions of rows and you run into OutOfMemory issues, consider increasing the `processor-batch-size` in the `<jboss-install>/standalone/configuration/standalone-teiid.xml` file to force the allocation of larger batches and table pages. If the processor batch size is increased and/or you are dealing with extremely wide result sets (several hundred columns), then the default setting of 8MB for the `max-storage-object-size` in the `<jboss-install>/standalone/configuration/standalone-teiid.xml` file may be too low. The sizing for `max-storage-object-size` is terms of serialized size, which will be much closer to the raw data size than the Java memory footprint estimation used for `max-reserved-kb`. `max-storage-object-size` should not be set too large relative to `memory-buffer-space` since it will reduce the performance of the memory buffer. The memory buffer supports only 1 concurrent writer for each `max-storage-object-size` of the `memory-buffer-space`.



#### Note

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

The `memory-buffer-space` setting controls the amount of on or off heap memory allocated as byte buffers for use by the Teiid 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 `max-reserve-kb`.

You can take advantage of the buffer manager memory buffer to access system memory without allocating it to the heap. Setting `memory-buffer-off-heap` to true in `<jboss-install>/standalone/configuration/standalone-teiid.xml` will allocate the Teiid memory buffer off heap. Depending on whether your installation is dedicated to Teiid and the amount of system memory available, this may be preferable to on-heap allocation. The primary benefit is additional memory usage for Teiid without additional garbage collection tuning. This becomes especially important in situations where more than 32GB of memory is desired for the VM. Note that when using off-heap allocation, the memory must still be available to the java process and that setting the value of `memory-buffer-space` too high may cause the VM to swap rather than reside in memory.

With large off-heap buffer sizes (greater than several gigabytes) you may also need to adjust VM settings. For Sun VMs the relevant VM settings are `MaxDirectMemorySize` and `UseLargePages`. For example adding:

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

to the VM process arguments would allow for an effective allocation of approximately an 11GB Teiid memory buffer (the `memory-buffer-space` setting) accounting for any additional direct memory that may be needed by the AS or applications running in the AS.

## 6.1.2. Disk Usage

For table page and result batches the buffer manager will use a limited number of files that are dedicated to a particular storage size. However, as mentioned in the installation, creation of Teiid 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 Teiid will use up to 50GB of disk space. This is tracked in terms of the number of bytes written by Teiid. For large data sets, you may need to increase the `'max-buffer-space'` setting in the `<jboss-install>/standalone/configuration/standalone-teiid.xml` file.

## 6.2. Threading

Socket threads are configured for each [transport](#). They 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 with work the query engine.

The query engine has several settings that determine its thread utilization. `max-threads` sets the total number of threads available for query engine work (processing plans, transaction control operations, processing source queries, etc.). You should consider increasing the maximum threads on systems with a large number of available processors and/or when it's common to issue non-transactional queries with that issue a large number of concurrent source requests. `max-active-plans`, which should always be smaller than `max-threads`, sets the number of the `max-threads` that should be used for user query processing. 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 the `max-active-plans`, then consider decreasing the processor/connector batch sizes to limit the base number of memory rows consumed by each plan. `thread-count-for-source-concurrency`, which should always be smaller than `max-threads`, sets the number of concurrently executing source queries per user request. Setting this value to 1 forces serial execution of all source queries by the processing thread. The default value is computed based upon  $2 * \text{max-threads} / \text{max-active-plans}$ . Using the respective default values, this means that 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.

### 6.3. Cache Tuning

Caching can be tuned for cached result (including user query results and procedure results) and prepared plans (including user and stored procedure plans). Even though it is possible to disable or otherwise severely constrain these caches, this would probably never be done in practice as it would lead to poor performance.

Cache statistics can be obtained through the Admin 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 Teiid.

While the result cache parameters control the cache result entries (max number, eviction, etc.), the result batches themselves are accessed through the [BufferManager](#). If the size of the result cache is increased, you may need to tune the BufferManager 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 Teiid. See the Developers Guide for further customization. Teiid stores compiled forms of update plans or trigger actions with the prepared plan, so it is recommended to leave the max-staleness of the prepared plan cache set to 0 so that metadata changes, for example disabling a trigger, may take effect immediately. The default staleness for result set caching is 60 seconds to improve efficiency with rapidly changing sources. Consider decreasing this value to make the result set cache more consistent with the underlying data. Even with a setting of 0 full transactional consistency is not guaranteed.

### 6.4. Socket Transports

Teiid separates the configuration of its socket transports for JDBC and ODBC. Typical installations will not need to adjust the default thread and buffer size settings. The default input output buffer sizes are set to 0, which will use the system default. Before adjusting this value keep in mind that each JDBC, ODBC will create a new socket connection. Setting these values to a large buffer size should only be done if the number of client is constrained. All JDBC/ODBC socket operations are non-blocking, so setting the number of max-threads higher than the maximum effective parallelism of the machine should not result in greater performance. The default value 0 for JDBC socket threads will set the max to the number of available processors.

### 6.5. LOBs

LOBs and XML documents are streamed from the Teiid Server to the Teiid 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 in the `<jboss-install>/standalone/configuration/standalone-teiid.xml` file. The default value is 100 KB. When dealing with extremely large LOBs, you may consider increasing this value 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.


## 6.6. Other Considerations

When using Teiid in a development environment, you may consider setting the `max-source-rows-allowed` property in the `<jboss-install>/standalone/configuration/standalone-teiid.xml` file to reasonably small level value (e.g. 10000) to prevent large amounts of data from being pulled from sources. Leaving the `exception-on-max-source-rows` set to true will alert the developer through an exception that an attempt was made to retrieve more than the specified number of rows.



# Teiid Admin Console

The Teiid Admin Console is a web based administrative and monitoring tool for Teiid. Teiid's Admin Console is built using the *Embedded JOPR* [<http://www.jboss.org/embjopr>] library and adds a additional plugin into the Embedded JOPR program already available in the *JBoss AS* [<http://www.jboss.org/jbossas>].



## 7.1. What can be monitored and/or configured?

Here are the steps to follow to install Teiid

1. *The Teiid Runtime Engine* (Data Services node in the tree)
2. *VDBs - Virtual databases*
  - a. *Models*
    - i. *Source*- these are physical sources
    - ii. *Multi-source* - these are multiple sourced models
    - iii. *Logical* - these are virtual sources
  - b. *Translator instances*- any Translator instances defined for use by this VDB
  - c. *Data Roles*- any data roles defined for use by this VDB
3. *Translators* - These are the extensions to supported datasources that come with Teiid out-of-the-box.



### Note

The creation/modification of the datasource is managed by the JBossAS plugin.

### 7.1.1. Configuration

1. Settings for the Data Services plugin
2. Runtime Engine properties
3. Buffer Service
4. Jdbc Socket configuration
5. Session Service

### 7.1.2. Metrics

1. Long Running Query count
2. Active Query count
3. Active Session count
4. Prepared Plan Cache Hit Ratio %
5. Prepared Plan Cache Size
6. Prepared Plan Cache # of Requests
7. ResultSet Cache Hit Ratio %
8. ResultSet Cache Size
9. ResultSet Cache # of Requests
10. Used Buffer Space

### 7.1.3. Control (Operations)

<listitem>

#### Data Services Engine

1. View Long Running Queries
2. View Current Sessions
3. Deploy a VDB via URL

- 4. Terminate Session
  - 5. View Current Requests
  - 6. Terminate requests
  - 7. View Current Transactions
  - 8. Terminate Transaction
- </listitem>  
</listitem>  
Virtual Database

- 1. View VDB Requests
  - 2. View VDB Sessions
  - 3. List Materialized View Info
  - 4. Refresh a Materialized View
  - 5. Clear Cache
- </listitem>

#### 7.1.4. Deploying the VDB

VDB archive files created in the Designer Tool or Dynamic VDBs can be deployed into Teiid server using the Admin Console.

1. Select the Virtual Database node in the Admin Console tree and click the Add New Resource button.
2. Select the VDB archive file from the file system and click continue.
3. The VDB will deploy if no fatal errors are found in the archive. The status of the VDB will be UP if no errors are found with the models in the VDB.
4. If there are model errors, the VDB will be deployed with a status of DOWN and the errors will be listed on the configuration tab of the VDB. VDBs that are not UP will be marked with a red X in the tree.

Only Model's "connection-jndi-name" can be modified using this tool by clicking on the "configuration" tab, all other properties are read-only.



# AdminShell

## 8.1. Introduction

The AdminShell tooling provides scripting based programming environments that enable user to access, monitor and control a Teiid Server. Both the command line and graphical console tools are built on functionality provide by the Groovy ( <http://groovy.codehaus.org/>) project. The AdminShell tools can be used in ad-hoc scripting mode or to run pre-defined scripts.

### AdminShell features:

1. fully functional programming environment with resource flow control and exception management. See [Groovy](http://groovy.codehaus.org/) [http://groovy.codehaus.org/] docs for the full power of the language.
2. quick administrative tool. The user can connect to a running Teiid Server and invoke any of the AdminAPI methods, such as "deploy" or "createDataSource", to control the Teiid System. Since this can be script driven, these tasks can be automated and re-run at a later time.
3. simplified data access tool. The user can connect to a VDB, issue any SQL commands, and view the results of the query via [Groovy Sql](http://groovy.codehaus.org/Database+features) [http://groovy.codehaus.org/Database+features] extensions.
4. migration tool. This can be used to develop scripts like moving the Virtual Databases (VDB), Connection Factories, and Configuration from one development environment to another. This will enable users to test and automate their migration scripts before production deployments.
5. testing tool. The JUnit ( <http://junit.org> [http://junit.org/] ) test framework is built in, see [Groovy Unit Tests](http://groovy.codehaus.org/Unit+Testing) [http://groovy.codehaus.org/Unit+Testing]. User can write regression tests for checking system health, or data integrity that can be used to validate a system functionality automatically instead of manual verification by QA personnel.

### 8.1.1. Download

AdminShell is distributed along with other Teiid downloads under "teiid-7.6-adminshell-dist.zip" name. Download and unzip this file to any directory. Once you have unzipped the file, in root directory you will find "adminshell" and "adminshell-console" executable scripts to launch the command line and graphical tools respectively.

Windows: Double click or execute "adminshell.cmd"

\*nix: Execute the "adminshell.sh" script

## 8.2. Getting Started

To learn the basics of [Groovy](http://groovy.codehaus.org/) [http://groovy.codehaus.org/] take a look at their documents and tutorials on their website.

Basic knowledge of the Java programming language and types is required in order to effectively design and develop scripts using the AdminShell. To learn Java language find learning resources at <http://java.sun.com> [http://java.sun.com/].

You can learn about the Teiid AdminAPI either using “adminHelp()” function or by using the JavaDocs.

AdminShell is a specialized version of Groovy which works in several different modes: interactive shell, graphical console, or script run mode. In interactive shell mode (launched via adminshell), the user can invoke connect to a live Teiid system and issue any ad-hoc commands to control the system. The interactive buffer can be used to develop a script and the interactive session input and output can be captured into a log file, more on this later in the document.

In graphical mode (lanced via adminshell-console), the user can develop and run scripts using a text editor that supports syntax highlighting.

In the script run mode, the user can execute/play back previously developed scripts. This mode especially useful to automate any testing or to perform any repeated configurations/migrations changes to a Teiid system.

### 8.2.1. Essential Rules

To use AdminShell successfully, there are some basic syntactical rules to keep in mind.

1. In interactive shell mode, most commands (as seen by the help command) are used to control shell behavior and are not general Groovy scripting constructs. Admin methods will typically be called using functional notation:

```
connectAsAdmin()
```

2. All commands and functions are case sensitive.
3. An ending semicolon is optional for Groovy statements.
4. If a function requires input parameter(s), they should be declared inside "(" and ")". A function may have more than one parameter. String parameters can be wrapped in double or single quotes. Example:

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

5. Other Java methods and classes can be used from your scripts, if the required Java class libraries are already in class path. You may place additional jars in the lib directory to have be automatically part of the class path. An example showing an import:

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

To execute the commands and arbitrary script in interactive mode you enter them first and press enter to execute, then enter the next line, so on.

To exit the tool in the interactive mode, first disconnect if you are connected to the Teiid system by executing "disconnect();" then type "exit". In the script mode, when execution of the script finishes the tool will exit automatically, however you still have to disconnect from Teiid system in the script.

Note: If SSL is turned on the Teiid server, you would need to adjust the connection URL and the client SSL settings as necessary (typically this will only be needed for 2-way SSL).

## 8.2.2. Help

The adminHelp() methods lists all the available administrative API methods in the AdminShell. Please note that none of the Groovy Shell commands or other available function calls will be shown in this list

```
adminHelp();
```

To get a specific definition about a method and it's required input parameters, use adminHelp("method")

```
adminHelp("deployVDB");

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

The sqlHelp() methods lists all Sql extension methods.

```
sqlHelp();
```

To get a specific definition about a method and its required input parameters, use `sqlHelp("method")`

### 8.2.3. Basic Commands

The list below contains some common commands used in AdminShell.

```
println "xxx"; // print something to console

adminHelp(); // shows all the available admin commands;

sql = connect(); // get an extended Groovy Sql connection using connection.properties file

sql.execute(<SQL>); // run any SQL command.

connectAsAdmin(); // connect as admin; no need have the vdb name. SQL commands will not
work under this connection

println getConnectionName(); // returns the current connection name

useConnection(<connection name>); // switches to using the given connection settings

disconnect(); // disconnects the current connection in the context
```

## 8.3. Executing a script file

To execute a script file "foo.groovy" in a directory "some/directory" in the interactive command line tool, execute as following

```
. some/directory/foo.groovy
```

"foo.groovy" is read into current context of the shell as if you typed in the whole document. If your script only contained method calls, you can explicitly invoke the call to execute.

Full execute syntax may also be used, and is required outside of the interactive command line tool:

```
evaluate("some/directory/foo.groovy" as File)
```

To execute the same file without entering interactive mode, run

```
./adminshell.sh . some/directory/foo.groovy
```

Parameters can be passed in as Java System properties. For example

```
./adminshell.sh -Dparam=value . some/directory/foo.groovy
```

Inside the script file, you can access these properties using `System.getProperty`

```
value = System.getProperty("param"); // will return "value"
```

## 8.4. Log File and Recorded Script file

During the interactive mode, input is recorded in a history file. This file can be accessed via the up arrow in the interactive shell.

User can also capture the commands entered during a interactive session to their own script file by using "startRecording" and "stopRecording" commands. For example,

```
record start directory/filename.txt  
<commands and script ..>  
record stop
```

All input and output between the start and stop are captured in the "directory/filename.txt" file. This gives the user an option to capture only certain portions of the interactive session and to later refine a script out of recorded file.

## 8.5. Default Connection Properties

The file "connection.properties" in the installation directory of the AdminShell defines the default connection properties with which user can connect to Teiid system. The following properties can be defined using this file

```
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 "connectionAsAdmin()" without any input parameters, will connect to the Teiid system using the properties defined in properties file. However, a user can always pass in parameters in the connect method to connect to a same or different server than one mentioned in the "connection.properties". Look all the all the different connect methods using the "adminHelp()" method.

Note that it is not secure to leave the passwords in clear text, as defined above. Please take necessary measures to secure the properties file, or do not use this feature and always pass in password interactively or some other secure way.

Note: At any given time user can be actively connected to more than one system or have more than one connection to same system. To manage the connections correctly each connection is created given a unique connection name. To learn more about this look at [Handling Multiple Connections](#).

### 8.6. Handling Multiple Connections

Using AdminShell, a user can actively manage more than one connection to a single or multiple Teiid systems. For example, two separate connections can be maintained, one to the development server and one to the integration server at the same time. This is possible because AdminShell supports a feature called named connections.

Every time a connection is made, the connection has an explicit or an implicitly assigned name.

If another connect command is executed then a new connection is made with a unique name and execution will be switched to use the new connection. The previous connection will be held as it is in its current state, and will not be closed.

You can use the following command to find out the current connection's name

```
name = getConnectionName();
```

Knowing the names of the connection that user is working with is important to switch the active connection to a previous connection. To switch the active connection, use the following command and supply the name of the connection to be used

```
useConnection("name");
```

If user supplies the same name as the active connection as they are currently participating in, then this operation will simply return with out any modifications. There is no limitation the number of simultaneous connections.

The following shows an example of using and switching between two connections.

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

//capture the connection name
conn1 = getConnectionName();

deploy("file.vdb")

// creates a second connection
connectAsAdmin();

conn2 = getConnectionName();

deploy("file.vdb")

// switch the connection to "conn1"
useConnection(conn1);

// close the connection in the "conn1"
disconnectAll();
```

## 8.7. Interactive Shell Nuances

The interactive shell uses a special shell interpreter and therefore has different behavior than just writing a script in Groovy. See the [Groovy Shell Documentation](http://groovy.codehaus.org/Groovy+Shell) [http://groovy.codehaus.org/Groovy+Shell] for more on its usage. Notable differences:

- Def statements do not define a variable in the context of the Shell, e.g. do not use `def x = 1`, use `x = 1`
- Shell commands (as seen through `help`) using the non-functional shell syntax are only available in the shell.

- Groovy classes using annotations cannot be parsed in the Shell.

---

# Appendix A. AdminShell Frequently Asked Questions

## A.1. Why won't the adminhelp command work in the Console tool?

The Console environment does not understand Shell commands (load, help, adminhelp, etc.), since they are not directly supported by Groovy. In the Console you should use the equivalent functional form / Groovy, e.g. instead of adminhelp, adminHelp()

## A.2. I have written a very useful script to do XYZ, I would like this to be part of the distribution?

Yes, we would love to hear from users. Please submit the script through the [Teiid JIRA](https://jira.jboss.org/jira/browse/TEIID) [https://jira.jboss.org/jira/browse/TEIID], and if this script popular, we will include the script in the scripts library in the following releases.

## A.3. What is different between "connectAsAdmin()" and "connect()"?

The connectAsAdmin methods create a contextual connection to the AdminAPI of the Teiid Server. The connect methods return an extension of the Groovy Sql object to be used for Sql calls to the Teiid Server.

## A.4. What does "getAdmin()" call do? Why do I need it?

"getAdmin()" returns this contextual connection object created when you executed "connectAsAdmin()" method. This object implements the interface "org.teiid.adminapi.Admin" and 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.

## A.5. Is IDE support available for writing the scripts?

The Admin Console tool is a light-weight IDE. Full IDE support is available for Groovy, but requires manual manipulation of the class path and script imports. See [using AdminShell methods in other environments](#).

## A.6. Is debugging support available?

The interactive shell and console do have built-in support for inspection of the current state. Performing line based debugging is beyond the scope of this document.



---

# Appendix B. Other Scripting Environments

The AdminShell methods (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 `<adminshell dist>/lib/teiid-7.6-client.jar` and `<adminshell dist>/lib/teiid-adminshell-7.6.jar` are in your class path. The snippet below show import statements that would work in Java, BeanShell, Groovy, etc.

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

Note that the provided shell and console executables automatically have the proper class path set and inject the proper imports into running scripts. If you wish to use scripts in a non-Teiid Groovy environment, you will need to manually add these imports and add the admin/client jars to the class path.



---

## Appendix C. System Properties

Some of Teiid's low-level behavior can be configured via system properties, rather than through configuration files. A typical place to set system properties for JBoss AS launches is in the <jboss-install>/bin/run.conf. A property setting has the format -Dproperty=value

- *org.teiid.allowNaNInfinity* - defaults to false. 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.
- *org.teiid.useValueCache* - defaults to false. Set to true to enable the canonical value cache. Value caching is used dynamically when buffer memory is consumed to reuse identical values and thus reduce the memory consumed by Teiid. 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.
- *org.teiid.ansiQuotedIdentifiers* - defaults to true. Set to false to emulate Teiid 6.x and prior behavior of treating double quoted values without leading identifier parts as string literals, which is not expected by the SQL specification.
- *org.teiid.subqueryUnnestDefault* - defaults to false. 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 by upon the costing information available.
- *org.teiid.ODBCPacketSize* - defaults to 307200. Target size in bytes of the ODBC results buffer. This is not a hard maximum, lobs and wide rows may use larger buffers.
- *org.teiid.iso8601Week* - defaults to false. Set to true to use ISO 8601 rules for week calculations regardless of the locale. When 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.
- *org.teiid.decimalAsDouble* - defaults to false. Set to true to parse exact fixed point literals, e.g. 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 releases earlier than 8.0.

