

JBoss Enterprise Application Platform

4.2.0.CP01

Getting Started Guide

ISBN: N/A

Publication date: Sep, 2007

JBoss Enterprise Application ...

The Getting Started Guide provides post-installation information about JBoss Enterprise Application Platform. Use this guide to familiarise yourself with the platform and the sample applications that demonstrate application development and deployment.

JBoss Enterprise Application Platform: Getting Started Guide

Copyright © 2007 Red Hat, Inc.

Copyright © 2007 by Red Hat, Inc. This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, V1.0 or later (the latest version is presently available at <http://www.opencontent.org/openpub/>).

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of the work or derivative of the work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

Red Hat and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

The GPG fingerprint of the security@redhat.com key is:

CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

1801 Varsity Drive
Raleigh, NC 27606-2072
USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701
PO Box 13588
Research Triangle Park, NC 27709
USA

About this book	vii
1. Introduction	1
1. Feedback	1
2. Other Manuals	1
2. The JBoss Server - A Quick Tour	3
1. Directory Structure	3
1.1. JBoss Top Level Directory Structure	3
1.2. JBOSS_DIST/jboss-as Directory Structure	4
2. Server Configurations	5
2.1. Server Configuration Directory Structure	6
2.2. The "default" Server Configuration File Set	7
2.3. The "all" Server Configuration File Set	13
2.4. EJB3 Services	13
2.5. Adding Your Own Configuration	14
3. Starting and Stopping the Server	14
3.1. Start the Server	14
3.2. Start the Server With Alternate Configuration	15
3.3. Using run.sh	16
3.4. Stopping the Server	16
3.5. Running as a Service under Microsoft Windows	17
4. The JMX Console	17
5. Hot-deployment of services in JBoss	19
6. Basic Configuration Issues	19
6.1. Core Services	19
6.2. Logging Service	20
6.3. Security Service	21
6.4. Additional Services	23
7. The Web Container - Tomcat	23
3. EJB3 Caveats in JBoss Enterprise Application Platform 4.2.0.CP01	25
1. Unimplemented features	25
2. Referencing EJB3 Session Beans from non-EJB3 Beans	25
4. About the Example Applications	27
1. Install Ant	27
5. Sample JSF-EJB3 Application	29
1. Data Model	29
2. JSF Web Pages	30
3. EJB3 Session Beans	34
4. Configuration and Packaging	36
4.1. Building The Application	36
4.2. Configuration Files	37
5. The Database	39
5.1. Creating the Database Schema	39
5.2. The HSQL Database Manager Tool	39
6. Deploying the Application	40
6. Using Seam	43
1. Data Model	43

- 2. JSF Web Pages - index.xhtml and create.xhtml44
- 3. Data Access using a Session Bean46
- 4. JSF Web Pages - todos.xhtml and edit.xhtml47
- 5. Xml Files49
- 6. Further Information50
- 7. Using other Databases51
 - 1. DataSource Configuration Files51
 - 2. Using MySQL as the Default DataSource51
 - 2.1. Creating a Database and User51
 - 2.2. Installing the JDBC Driver and Deploying the DataSource52
 - 2.3. Testing the MySQL DataSource53
 - 3. Setting up an XADataSource with Oracle 9i53
 - 3.1. Padding Xid Values for Oracle Compatibility54
 - 3.2. Installing the JDBC Driver and Deploying the DataSource54
 - 3.3. Testing the Oracle DataSource55
- A. Further Information Sources57

About this book

The goal of this book is to give you an overview of JBoss Enterprise Application Platform, and demonstrate some of its features and capability to provide a rapid application development and deployment environment for Enterprise Java Applications. At the time of writing, the latest release is version 4.2.0.CP01. You should use this version or later with the example applications. The example applications described in this book illustrate the development and deployment of Enterprise Java applications in JBoss Enterprise Application Platform. While the book is not intended to teach you Enterprise Java development, we will be covering the subject from quite a basic standpoint, so it will still be useful if you are new to Enterprise Java Application Development.

We will provide you a quick tour of the JBoss Directory Structure, basic Server Configuration, key configuration files and the JMX and Web Consoles. As we move on to the example applications, you will see JBoss Enterprise Application Platform in action and get some exposure in simple configuration and deployment issues. We will introduce the Seam framework and demonstrate the significant difference that Seam makes to application development.

Of course, that barely scratches the surface of what you can do with JBoss Enterprise Application Platform. Once you feel comfortable with the information here, the *JBoss Enterprise Application Platform: Server Configuration Guide* can take you through the rest of the way to total mastery of JBoss Enterprise Application Platform.

Introduction

JBoss Enterprise Application Platform is easy to install and you can have it running in a few easy steps. Refer to the *JBoss Enterprise Application Platform: Installation Guide* for information on pre-requisites for installation and the detailed installation steps.

Once you have JBoss Enterprise Application Platform installed, use this guide to familiarize yourself with its layout and the example applications that demonstrate application development and deployment.

1. Feedback

If you spot a typo in the *JBoss Enterprise Application Platform: Getting Started Guide*, or if you have thought of a way to make this manual better, we would love to hear from you! Submit a report in *Bugzilla* [<http://bugzilla.redhat.com/bugzilla/>] against the Product: JBoss Enterprise Application Platform, Version: 4.2.0, Component: Getting_Started_Guide. If you have a suggestion for improving the documentation, try to be as specific as possible. If you have found an error, include the section number and some of the surrounding text so we can find it easily.

2. Other Manuals

If you are looking for detailed product information refer to the manuals available online at <http://www.redhat.com/docs/manuals/jboss>.

The JBoss Server - A Quick Tour

Let's explore the JBoss Enterprise Application Platform directory structure and help you understand how the installation is laid out and what goes where. It's worth familiarizing yourself with the layout, locations of the key configuration files, log files, deployment and so on. It will help you understand the JBoss service architecture so that you'll be able to find your way around when it comes to deploying your own applications.

1. Directory Structure

1.1. JBoss Top Level Directory Structure

Installing JBoss Enterprise Application Platform creates a top level directory, which will be named `jboss-eap-4.2` if you used the zip installation method, and will be named according to your specification if you used the GUI installer. Throughout this guide we refer to this top-level directory as the `JBOSS_DIST` directory. There are four sub-directories immediately below this:

- **doc**: contains the product documentation.
- **jboss-as**: contains sub directories that contain server start scripts, JARs, server configuration sets and working directories. You need to know your way around the distribution layout to locate JARs for compiling code, updating configurations, deploying your code, etc.
- **seam**: contains the files for Hibernate and the JBoss Seam Framework.
- **Uninstaller**: contains the uninstaller program `uninstaller.jar`.

Below is the layout of the installation directory of JBoss Enterprise Application Platform. In the figure, the `default` server configuration file set is shown expanded. In a clean installation, within the `server/default` directory only the `conf`, `deploy`, and `lib` directories exist. The `data`, `log`, `tmp` and `work` sub-directories are created by JBoss and won't exist until you've run the server at least once. [Section 3, "Starting and Stopping the Server"](#) will teach you to run the server.

```
jboss-eap-4.2           // jboss.home_url
  |+ doc
  |+ jboss-as
      |+ bin
      |+ client
      |+ docs
      |+ icons
      |+ lib             // jboss.lib.url
      |+ scripts
      |+ server
          |+ all         // jboss.server.name
          |+ default     // jboss.server.home.url
              |+ conf    // jboss.server.config.url
              |+ deploy
              |+ lib     // jboss.server.lib.url
```

```

|+ data
|+ log
|+ tmp
|+ work
|+ minimal
|+ production
|+ seam
|+ Uninstaller // jboss.uninstaller.url

```

Several of the locations may be overridden. For these locations, the `org.jboss.system.server.ServerConfig` interface constant and its corresponding system property string are shown in the figure. The names ending in `URL` correspond to locations that can be specified using a URL to access remote locations, for example, HTTP URLs against a web server.

1.2. JBOSS_DIST/jboss-as Directory Structure

The table below illustrates the contents of the `jboss-as` directory.

Directory	Description
<code>bin</code>	Contains startup, shutdown and other system-specific scripts. Basically all the entry point JARs and start scripts included with the JBoss distribution are located in the <code>bin</code> directory.
<code>client</code>	Stores configuration files and JAR files that may be used by a Java client application (running outside JBoss) or an external web container. You can select archives as required or use <code>jbossall-client.jar</code> .
<code>docs</code>	Contains the XML DTDs used in JBoss for reference (these are also a useful source of documentation on JBoss configuration specifics). There are also example JCA (Java Connector Architecture) configuration files for setting up datasources for different databases (such as MySQL, Oracle, Postgres).
<code>lib</code>	Contains startup JARs used by JBoss. Do not place your own JAR files in this directory.
<code>server</code>	Contains the JBoss server configuration sets. Each of the subdirectories in here is a different server configuration. JBoss ships with <code>minimal</code> , <code>default</code> , <code>production</code> , and <code>all</code> configuration sets. The subdirectories and key configuration files contained in the <code>default</code>

Directory	Description
	configuration set are discussed in more detail in subsequent sections.

Table 2.1. Contents of `JBOSS_DIST/jboss-as` directory

2. Server Configurations

Fundamentally, the JBoss architecture consists of the JMX MBean server, the microkernel, and a set of pluggable component services - the MBeans. This makes it easy to assemble different configurations and gives you the flexibility to tailor them to meet your requirements.

You don't have to run a large, monolithic server all the time; you can remove the components you don't need (which can also reduce the server startup time considerably) and you can also integrate additional services into JBoss by writing your own MBeans. You certainly don't need to do this to be able to run standard J2EE applications though. Everything you need is already there.

You don't need a detailed understanding of JMX to use JBoss, but it's worth keeping a picture of this basic architecture in mind as it is central to the way JBoss works.

The JBoss Enterprise Application Platform ships with four different server configurations. Within the `JBOSS_DIST/jboss-as/server` directory, you will find four subdirectories: `minimal`, `default`, `production` and `all` - one for each server configuration. Each of these configurations provide a different set of services. The `production` configuration is the one used if you don't specify another one when starting up the server.

minimal

has a minimal configuration—the bare minimum services required to start JBoss. It starts the logging service, a JNDI server and a URL deployment scanner to find new deployments. This is what you would use if you want to use JMX/JBoss to start your own services without any other J2EE technologies. This is just the bare server. There is no web container, no EJB or JMS support. This is not a J2EE 1.4 compatible configuration.

default

is a base J2EE 1.4 server profile containing a default set of services. It has the most frequently used services required to deploy a J2EE application. It does not include the JAXR service, the IIOP service, or any of the clustering services. Please note that although this configuration is called "default", the actual default configuration for the server is the "production" configuration.

all

on the other hand has all the services configured to launch every single component. This is a full J2EE 1.4 server profile with enterprise extensions such as Clustering and RMI/IIOP.

production

is based on the "all" profile, tuned for production; with log verbosity reduced, deployment scanning every 60 seconds, and memory usage tuned to accommodate production deployment requirements, among other things. This is the configuration that will be used by the server when it is started, if no other configuration is specified.

If you want to know which services are configured in each of these instances, look at the `jboss-service.xml` file in the `JBOSS_DIST/jboss-as/server/<instance-name>/conf/` directory and also the configuration files in the `JBOSS_DIST/jboss-as/server/<instance-name>/deploy` directory.

```
[vsr]$ls server/default/conf
jbossjta-properties.xml  jndi.properties  standardjbosscomp-jdbc.xml
jboss-log4j.xml         login-config.xml  standardjboss.xml
jboss-minimal.xml      props             xmdesc
jboss-service.xml      standardjaws.xml
```



Note

The **production** configuration is the one used if you don't specify another one when starting up the server.

To start the server using an alternate configuration refer to [Section 3.2, "Start the Server With Alternate Configuration"](#).

2.1. Server Configuration Directory Structure

The directory server configuration you're using, is effectively the server root while JBoss is running. It contains all the code and configuration information for the services provided by the particular server configuration. It's where the log output goes, and it's where you deploy your applications. [Table 2.2, "Server Configuration Directory Structure"](#) shows the directories inside the server configuration directory (`JBOSS_DIST/jboss-as/server/<instance-name>`) and their functions.

Directory	Description
conf	The <code>conf</code> directory contains the <code>jboss-service.xml</code> bootstrap descriptor file for a given server configuration. This defines the core services that are fixed for the lifetime of the server.
data	The <code>data</code> directory is available for use by services that want to store content in the file system. It holds persistent data for services intended to survive a server restart. Several JBoss services, such as the embedded Hypersonic database instance, store data here.

Directory	Description
deploy	The <code>deploy</code> directory contains the hot-deployable services (those which can be added to or removed from the running server). It also contains applications for the current server configuration. You deploy your application code by placing application packages (JAR, WAR and EAR files) in the <code>deploy</code> directory. The directory is constantly scanned for updates, and any modified components will be re-deployed automatically. This may be overridden through the <code>URLDeploymentScanner</code> URLs attribute.
lib	This directory contains JAR files (Java libraries that should not be hot deployed) needed by this server configuration. You can add required library files here for JDBC drivers etc. All JARs in this directory are loaded into the shared classpath at startup.
log	This is where the log files are written. JBoss uses the Jakarta <code>log4j</code> package for logging and you can also use it directly in your own applications from within the server. This may be overridden through the <code>conf/log4j.xml</code> configuration file.
tmp	The <code>tmp</code> directory is used for temporary storage by JBoss services. The deployer, for example, expands application archives in this directory.
work	This directory is used by Tomcat for compilation of JSPs.

Table 2.2. Server Configuration Directory Structure

2.2. The "default" Server Configuration File Set

The "default" server configuration file set is located in the `JBOSS_DIST/jboss-as/server/default` directory. Let's take a look at the contents of the default server configuration file set:

```

jboss-eap-4.2           // jboss.home_url
  |+ doc
  |+ jboss-as
    |+ bin
    |+ client
    |+ docs
    |+ icons
    |+ lib               // jboss.lib.url
    |+ scripts
    |+ server
      |+ all             // jboss.server.name
      |+ default        // jboss.server.home.url
      |+ conf           // jboss.server.config.url
      |+ props
      |+ xmdesc
      - jbossjta-properties.xml
      - jboss-minimal.xml
      - jndi.properties

```

```

- standardjboss.xml
- jboss-log4j.xml
- jboss-service.xml
- login-config.xml
- standardjbosscomp-jdbc.xml
|+ deploy
|+ ejb3.deployer
|+ http-invoker.sar
|+ jboss-aop-jdk50.deployer
|+ jboss-bean.deployer
|+ jboss-web.deployer
|+ jbossws.sar
|+ jms
|+ jmx-console.war
|+ management
|+ uuid-key-generator.sar
- bsh-deployer.xml
- cache-invalidation-service.xml
- client-deployer-service.xml
- ear-deployer.xml
- ejb3-interceptors-aop.xml
- ejb-deployer.xml
- hsqldb-ds.xml
- jboss-ha-local-jdbc.rar
- jboss-ha-xa-jdbc.rar
- jbossjca-service.xml
- jboss-local-jdbc.rar
- jboss-xa-jdbc.rar
- jmx-invoker-service.xml
- jsr88-service.xml
- mail-service.xml
- monitoring-service.xml
- properties-service.xml
- quartz-ra.rar
- schedule-manager-service.xml
- scheduler-service.xml
- sqlexception-service.xml
|+ lib // jboss.server.lib.url
|+ minimal
|+ production
|+ seam
|+ Uninstaller // jboss.uninstaller.url

```

2.2.1. Contents of "conf" directory

The files in the `conf` directory are explained in the following table.

File	Description
<code>jboss-minimal.xml</code>	This is a minimalist example of the <code>jboss-service.xml</code> configuration file. (This is the <code>jboss-service.xml</code> file used in the minimal configuration file set)

File	Description
<code>jboss-service.xml</code>	<code>jboss-service.xml</code> defines the core services and their configurations.
<code>jndi.properties</code>	The <code>jndi.properties</code> file specifies the JNDI <code>InitialContext</code> properties that are used within the JBoss server when an <code>InitialContext</code> is created using the no-arg constructor.
<code>jboss-log4j.xml</code>	This file configures the Apache log4j framework category priorities and appenders used by the JBoss server code.
<code>login-config.xml</code>	This file contains sample server side authentication configurations that are applicable when using JAAS based security.
<code>props/ *</code>	The <code>props</code> directory contains the users and roles property files for the <code>jmx-console</code> .
<code>standardjaws.xml</code>	This file provides the default configuration for the legacy EJB 1.1 CMP engine.
<code>standardjboss.xml</code>	This file provides the default container configurations.
<code>standardjbosscmp-jdbc.xml</code>	This file provides a default configuration file for the JBoss CMP engine.
<code>xmdesc/ *-mbean.xml</code>	The <code>xmdesc</code> directory contains XMLElement descriptors for several services configured in the <code>jboss-service.xml</code> file.

Table 2.3. Contents of "conf" directory

2.2.2. Contents of "deploy" directory

The files in the `deploy` directory are explained in the following table.

File	Description
<code>bsh-deployer.xml</code>	This file configures the bean shell deployer, which deploys bean shell scripts as JBoss services.
<code>cache-invalidation-service.xml</code>	This is a service that allows for custom invalidation of the EJB caches via JMS notifications. It is disabled by default.
<code>client-deployer-service.xml</code>	This is a service that provides support for J2EE application clients. It manages the

File	Description
	java:comp/env enterprise naming context for client applications based on the application-client.xml descriptor.
ear-deployer.xml	The EAR deployer is the service responsible for deploying J2EE EAR files.
ejb-deployer.xml	The EJB deployer is the service responsible for deploying J2EE EJB JAR files.
hsqldb-ds.xml	hsqldb-ds.xml configures the Hypersonic embedded database service configuration file. It sets up the embedded database and related connection factories.
http-invoker.sar	http-invoker.sar contains the detached invoker that supports RMI over HTTP. It also contains the proxy bindings for accessing JNDI over HTTP.
jboss-aop-jdk50.deployer	This service configures the AspectManagerService and deploys JBoss AOP applications.
jboss-bean.deployer	jboss-bean.deployer provides the JBoss microcontainer, which deploys POJO services wrapped in .beans files.
jboss-ha-local-jdbc.rar	jboss-ha-local-jdbc.rar is an experimental version of jboss-local-jdbc.rar that supports datasource failover.
jboss-ha-xa-jdbc.rar	jboss-ha-xa-jdbc.rar is an experimental version of jboss-xa-jdbc.rar that supports datasource failover.
jboss-local-jdbc.rar	jboss-local-jdbc.rar is a JCA resource adaptor that implements the JCA ManagedConnectionFactory interface for JDBC drivers that support the DataSource interface but not JCA.
jboss-xa-jdbc.rar	jboss-xa-jdbc.rar is a JCA resource adaptor that implements the JCA ManagedConnectionFactory interface for JDBC drivers that support the XADataSource interface.
jbossjca-service.xml	jbossjca-service.xml is the application server implementation of the JCA specification. It provides the connection

File	Description
	management facilities for integrating resource adaptors into the JBoss server.
<code>jboss-web.deployer</code>	The <code>jboss-web.deployer</code> directory provides the Tomcat servlet engine.
<code>jbossws.sar</code>	<code>jbossws.sar</code> provides J2EE web services support.
<code>jms/hsqldb-jdbc-state-service.xml</code>	<code>hsqldb-jdbc-state-service.xml</code> provides JMS state management using Hypersonic.
<code>jms/hsqldb-jdbc2-service.xml</code>	<code>hsqldb-jdbc2-service.xml</code> configures JMS persistence and caching using Hypersonic. It also contains the <code>DestinationManager</code> MBean, which is the core service for the JMS implementation.
<code>jms/jbossmq-destinations-service.xml</code>	<code>jbossmq-destinations-service.xml</code> configures a number of JMS queues and topics used by the JMS unit tests.
<code>jms/jbossmq-httpil.sar</code>	<code>jbossmq-httpil.sar</code> provides a JMS invocation layer that allows the use of JMS over HTTP.
<code>jms/jbossmq-service.xml</code>	The <code>jbossmq-service.xml</code> file configures the core JBossMQ JMS service.
<code>jms/jms-ds.xml</code>	The <code>jms-ds.xml</code> file configures the JBossMQ JMS provider for use with the <code>jms-ra.rar</code> JCA resource adaptor.
<code>jms/jms-ra.rar</code>	<code>jms-ra.rar</code> is a JCA resource adaptor that implements the JCA <code>ManagedConnectionFactory</code> interface for JMS connection factories.
<code>jms/jvm-il-service.xml</code>	<code>jvm-il-service.xml</code> configures the in-JVM JMS transport invocation layer.
<code>jms/ui12-service.xml</code>	<code>ui12-service.xml</code> configures the JMS version 2 unified invocation layer. Its a fast and reliable custom socket based transport that should be used for messaging between JVMs.
<code>jmx-console.war</code>	The <code>jmx-console.war</code> directory provides the JMX Console. The JMX Console provides a simple web interface for managing the MBean server.
<code>jmx-invoker-service.sar</code>	<code>jmx-invoker-service.sar</code> is an unpacked MBean service archive that exposes a subset

File	Description
	of the JMX <code>MBeanServer</code> interface methods as an RMI interface to enable remote access to the JMX core functionality. This is similar to the legacy <code>jmx-rmi-adaptor.sar</code> , with the difference that the transport is handled by the detached invoker architecture.
<code>jsr-88-service.xml</code>	<code>jsr-88-service.xml</code> provides the JSR 88 remote deployment service.
<code>mail-ra.rar</code>	<code>mail-ra.rar</code> is a resource adaptor that provides a JavaMail connector.
<code>mail-service.xml</code>	The <code>mail-service.xml</code> file is an MBean service descriptor that provides JavaMail sessions for use inside the JBoss server.
<code>management/console-mgr.sar</code>	<code>console-mgr.sar</code> provides the Web Console. It is a web application/applet that provides a richer view of the JMX server management data than the JMX console. You may view the console using the URL http://localhost:8080/web-console/ .
<code>monitoring-service.xml</code>	The <code>monitoring-service.xml</code> file configures alert monitors like the console listener and email listener used by JMX notifications.
<code>properties-service.xml</code>	The <code>properties-service.xml</code> file is an MBean service descriptor that allows for customization of the JavaBeans <code>PropertyEditors</code> as well as the definition of system properties.
<code>scheduler-service.xml</code>	The <code>scheduler-service.xml</code> and <code>schedule-manager-service.xml</code> files are MBean service descriptors that provide a scheduling type of service.
<code>sqlexception-service.xml</code>	The <code>sqlexception-service.xml</code> file is an MBean service descriptor for the handling of vendor specific <code>SQLExceptions</code> .
<code>uuid-key-generator.sar</code>	The <code>uuid-key-generator.sar</code> service provides a UUID-based key generation facility.

Table 2.4. Contents of "deploy" directory

2.3. The "all" Server Configuration File Set

The "all" server configuration file set is located in the `JBOSS_DIST/jboss-as/server/all` directory. In addition to the services in the "default" set, the all configuration contains several other services in the `conf/` directory as shown below.

File	Description
<code>cluster-service.xml</code>	This service configures clustering communication for most clustered services in JBoss.
<code>deploy-hasingleton-service.xml</code>	This provides the HA singleton service, allowing JBoss to manage services that must be active on only one node of a cluster.
<code>deploy.last/farm-service.xml</code>	<code>farm-service.xml</code> provides the farm service, which allows for cluster-wide deployment and undeployment of services.
<code>httpha-invoker.sar</code>	This service provides HTTP tunneling support for clustered environments.
<code>iiop-service.xml</code>	This provides IIOP invocation support.
<code>juddi-service.sar</code>	This service provides UDDI lookup services.
<code>snmp-adaptor.sar</code>	This is a JMX to SNMP adaptor. It allows for the mapping of JMX notifications onto SNMP traps.
<code>tc5-cluster.sar</code>	Provides AOP support for field-level HTTP session replication.

Table 2.5. Additional Services in "conf" directory for "all" configuration

2.4. EJB3 Services

The following table explains the files providing ejb3 services.

File	Description
<code>ejb3-interceptors-aop.xml</code>	This service provides the AOP interceptor stack configurations for EJB3 bean types.
<code>ejb3.deployer</code>	This service deploys EJB3 applications into JBoss.
<code>jboss-aop-jdk50.deployer</code>	This is a Java 5 version of the AOP deployer. The AOP deployer configures the <code>AspectManagerService</code> and deploys JBoss AOP applications.

File	Description
jbossws.sar	This provides Java EE 5 web services support.

Table 2.6. EJB3 Services

Finally, in the EJB3 "all" configuration there are two additional services.

File	Description
ejb3-clustered-sfsbcache-service.xml	This provides replication and failover for EJB3 stateful session beans.
ejb3-entity-cache-service.xml	This provides a clustered cache for EJB3 entity beans.

Table 2.7. Additional Services in EJB3 "all" Configuration

2.5. Adding Your Own Configuration

You can add your own configurations too. The best way to do this is to copy an existing one that is closest to your needs and modify the contents. For example, if you weren't interested in using messaging, you could copy the `production` directory, renaming it as `myconfig`, remove the `jms` subdirectory and then start JBoss with the new configuration.

```
run -c myconfig
```

3. Starting and Stopping the Server

3.1. Start the Server

Move to `JBOSS_DIST/jboss-as/bin` directory and execute the `run.bat` (for Windows) or `run.sh` (for Linux) script, as appropriate for your operating system. Your output should look like the following (accounting for installation directory differences) and contain no error or exception messages:

```
[jwulf@thinkpad bin]$ ./run.sh
=====
JBoss Bootstrap Environment

JBOSS_HOME: /home/jwulf/jboss-eap-4.2/jboss-as

JAVA: java
```

```

    JAVA_OPTS: -Dprogram.name=run.sh -server -Xms1503m -Xmx1503m
-Dsun.rmi.dgc.client.
gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000
-Djava.net.preferIPv4Stack=true

    CLASSPATH: /home/jwulf/jboss-eap-4.2/jboss-as/bin/run.jar

=====

13:11:46,215 INFO [Server] Starting JBoss (MX MicroKernel)...
13:11:46,217 INFO [Server] Release ID: JBoss [EAP] 4.2.0.GA (build:
SVNTag=JBoss_4_2_0_GA date=200706111042)
13:11:46,218 INFO [Server] Home Dir: /home/jwulf/jboss-eap-4.2/jboss-as
13:11:46,219 INFO [Server] Home URL:
file:/home/jwulf/jboss-eap-4.2/jboss-as/
13:11:46,220 INFO [Server] Patch URL: null
13:11:46,225 INFO [Server] Server Name: production
13:11:46,225 INFO [Server] Server Home Dir:
/home/jwulf/jboss-eap-4.2/jboss-as/server/production
13:11:46,225 INFO [Server] Server Home URL:
file:/home/jwulf/jboss-eap-4.2/jboss-as/server/production/
13:11:46,225 INFO [Server] Server Log Dir:
/home/jwulf/jboss-eap-4.2/jboss-as/server/production/log
13:11:46,226 INFO [Server] Server Temp Dir:
/home/jwulf/jboss-eap-4.2/jboss-as/server/production/tmp
13:11:46,226 INFO [Server] Root Deployment Filename: jboss-service.xml
13:11:47,071 INFO [ServerInfo] Java version: 1.5.0_11,Sun Microsystems Inc.
13:11:47,071 INFO [ServerInfo] Java VM:
Java HotSpot(TM) Server VM 1.5.0_11-b03,Sun Microsystems Inc.
13:11:47,072 INFO [ServerInfo] OS-System: Linux 2.6.21-1.3228.rhel5,i386
13:11:48,558 INFO [Server] Core system initialized
13:11:56,934 INFO [WebService] Using RMI server codebase:
http://127.0.0.1:8083/
13:11:56,940 INFO [Log4jService$URLWatchTimerTask]
Configuring from URL: resource:jboss-log4j.xml

```



Note

Note that there is no "Server Started" message shown at the console when the server is started using the `production` profile, which is the default profile used when no other is specified. This message may be observed in the `server.log` file located in the `server/production/log` subdirectory.

3.2. Start the Server With Alternate Configuration

Using `run.sh` without any arguments starts the server using the `production` server configuration file set. To start with an alternate configuration file set, pass the name of the server configuration file set [same as the name of the server configuration directory under

`JBOSS_DIST/jboss-as/server]` that you want to use, as the value to the `-c` command line option. For example, to start with the `minimal` configuration file set you should specify:

```
[bin]$ ./run.sh -c minimal
...
...
...
15:05:40,301 INFO [Server] JBoss (MX MicroKernel) [4.2.0.GA (build:
SVNTag=JBoss_4_2_0_GA date=200706111042)] Started in 5s:75ms
```

3.3. Using run.sh

The `run` script supports the following options:

```
usage: run.sh [options]
-h, --help                Show help message
-V, --version            Show version information
--                        Stop processing options
-D<name>[=<value>]      Set a system property
-d, --bootdir=<dir>      Set the boot patch directory; Must be absolute
or url
-p, --patchdir=<dir>     Set the patch directory; Must be absolute or
url
-n, --netboot=<url>      Boot from net with the given url as base
-c, --configuration=<name> Set the server configuration name
-B, --bootlib=<filename> Add an extra library to the front
bootclasspath
-L, --library=<filename> Add an extra library to the loaders classpath
-C, --classpath=<url>    Add an extra url to the loaders classpath
-P, --properties=<url>   Load system properties from the given url
-b, --host=<host or ip>  Bind address for all JBoss services
-g, --partition=<name>   HA Partition name (default=DefaultDomain)
-u, --udp=<ip>           UDP multicast address
-l, --log=<log4j|jdk>    Specify the logger plugin type
```

3.4. Stopping the Server

To shutdown the server, you simply issue a `Ctrl-C` sequence in the console in which JBoss was started. Alternatively, you can use the `shutdown.sh` command.

```
[bin]$ ./shutdown.sh -S
```

The `shutdown` script supports the following options:

```
A JMX client to shutdown (exit or halt) a remote JBoss server.

usage: shutdown [options] <operation>

options:
-h, --help                Show this help message (default)
```

```

-D<name>[=<value>]      Set a system property
--                      Stop processing options
-s, --server=<url>      Specify the JNDI URL of the remote server
-n, --serverName=<url> Specify the JMX name of the ServerImpl
-a, --adapter=<name>   Specify JNDI name of the MBeanServerConnection to
use
-u, --user=<name>       Specify the username for authentication
-p, --password=<name>  Specify the password for authentication

operations:
-S, --shutdown          Shutdown the server
-e, --exit=<code>       Force the VM to exit with a status code
-H, --halt=<code>      Force the VM to halt with a status code

```

Using the shutdown command requires a server configuration that contains the `jmx-invoker-service.xml` service. Hence you cannot use the shutdown command with the minimal configuration.

3.5. Running as a Service under Microsoft Windows

You can configure the server to run as a service under Microsoft Windows, and configure it to start automatically if desired.

Download the `JavaService` package from <http://forge.objectweb.org/projects/javaservice/>.

Unzip the package and use the `JBossInstall.bat` file to install the JBoss service. You must set the `JAVA_HOME` and `JBOSS_HOME` environment variables to point to the `jdk` and `jboss-as` directories before running `JBossInstall.bat`. Run `JBossInstall.bat` with the following syntax:

```
JBossInstall.bat <depends> [-auto | -manual]
```

Where `<depends>` is the name of any service that the JBoss AS server depends on, such as the `mysql` database service.

Once the service is installed the server can be started by using the command `net start JBoss`, and stopped with the command `net stop JBoss`.

Please refer to the documentation included in the `JavaService` package for further information.

4. The JMX Console

When the JBoss Server is running, you can get a live view of the server by going to the JMX console application at <http://localhost:8080/jmx-console>. You should see something similar to [Figure 2.1, "View of the JMX Management Console Web Application"](#).

The JMX Console is the JBoss Management Console which provides a raw view of the JMX MBeans which make up the server. They can provide a lot of information about the running

server and allow you to modify its configuration, start and stop components and so on.

For example, find the `service=JNDIView` link and click on it. This particular MBean provides a service to allow you to view the structure of the JNDI namespaces within the server. Now find the operation called `list` near the bottom of the MBean view page and click the `invoke` button. The operation returns a view of the current names bound into the JNDI tree, which is very useful when you start deploying your own applications and want to know why you can't resolve a particular EJB name.

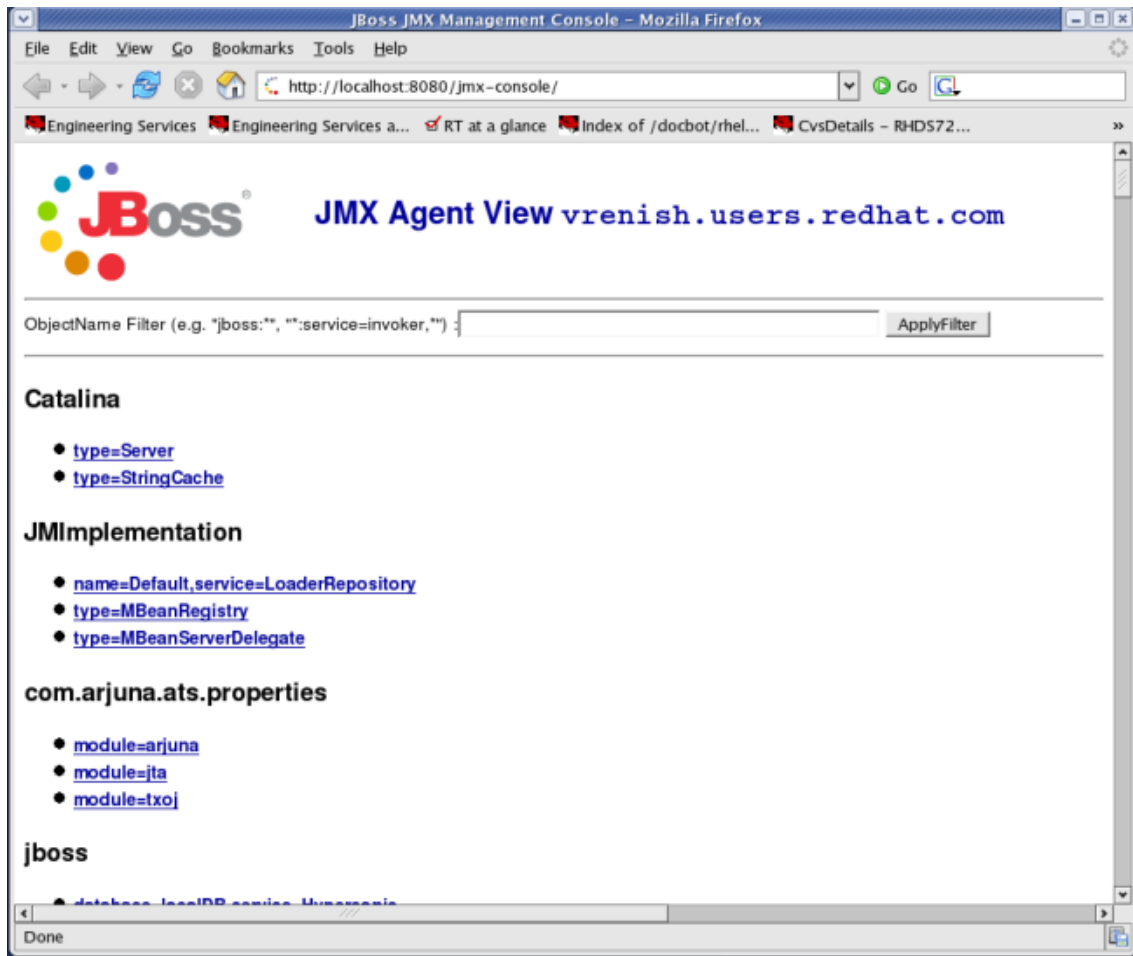


Figure 2.1. View of the JMX Management Console Web Application

Look at some of the other MBeans and their listed operations; try changing some of the configuration attributes and see what happens. With a very few exceptions, none of the changes made through the console are persistent. The original configuration will be reloaded when you restart JBoss, so you can experiment freely without doing any permanent damage.



Note

If you installed JBoss using the graphical installer, the JMX Console will prompt you for a username and password before you can access it. If you installed using other modes, you can still configure JMX Security manually. We will show you how to secure your console in [Section 6.3, "Security Service"](#).

5. Hot-deployment of services in JBoss

Hot-deployable services are those which can be added to or removed from the running server. These are placed in the `JBOSS_DIST/jboss-as/server/<instance-name>/deploy` directory. Let's have a look at a practical example of hot-deployment of services in JBoss before we go on to look at server configuration issues in more detail.

Start JBoss if it isn't already running and take a look at the `server/production/deploy` directory. Remove the `mail-service.xml` file and watch the output from the server:

```
13:10:05,235 INFO [MailService] Mail service 'java:/Mail' removed from JNDI
```

Then replace the file and watch JBoss re-install the service:

```
13:58:54,331 INFO [MailService] Mail Service bound to java:/Mail
```

This is hot-deployment in action.

6. Basic Configuration Issues

Now that we have examined the JBoss server, we will take a look at some of the main configuration files and what they are used for. All paths are relative to the server configuration directory (`server/production`, for example).

6.1. Core Services

The core services specified in the `conf/jboss-service.xml` file are started first when the server starts up. If you have a look at this file in an editor you will see MBeans for various services including logging, security, JNDI, JNDIView etc. Try commenting out the entry for the `JNDIView` service.

Note that because the mbeans definition had nested comments, we had to comment out the mbean in two sections, leaving the original comment as it was.

```
<!-- Section 1 commented out
<mbean code="org.jboss.naming.JNDIView"
      name="jboss:service=JNDIView"
```

```
    xmbean-dd="resource:xmdesc/JNDIView-xmbean.xml ">
-->
    <!-- The HANamingService service name -->
<!-- Section two commented out
    <attribute
name="HANamingService">jboss:service=HAJNDI</attribute></mbean>
-->
```

If you then restart JBoss, you will see that the `JNDIView` service no longer appears in the JMX Management Console (JMX Console) listing. In practice, you should rarely, if ever, need to modify this file, though there is nothing to stop you adding extra MBean entries in here if you want to. The alternative is to use a separate file in the `deploy` directory, which allows your service to be hot deployable.

6.2. Logging Service

In JBoss `log4j` is used for logging. If you are not familiar with the `log4j` package and would like to use it in your applications, you can read more about it at the Jakarta web site (<http://jakarta.apache.org/log4j/>).

Logging is controlled from a central `conf/log4j.xml` file. This file defines a set of appenders specifying the log files, what categories of messages should go there, the message format and the level of filtering. By default, JBoss produces output to both the console and a log file (`server.log` in the `log` directory).

There are 5 basic log levels used: `DEBUG`, `INFO`, `WARN`, `ERROR` and `FATAL`. The logging threshold on the console is `INFO`, which means that you will see informational messages, warning messages and error messages on the console but not general debug messages. In contrast, there is no threshold set for the `server.log` file, so all generated logging messages will be logged there.

If things are going wrong and there doesn't seem to be any useful information in the console, always check the `server.log` file to see if there are any debug messages which might help you to track down the problem. However, be aware that just because the logging threshold allows debug messages to be displayed, that doesn't mean that all of JBoss will produce detailed debug information for the log file. You will also have to boost the logging limits set for individual categories. Take the following category for example.

```
<!-- Limit JBoss categories to INFO -->
<category name="org.jboss">
    <priority value="INFO"/>
</category>
```

This limits the level of logging to `INFO` for all JBoss classes, apart from those which have more specific overrides provided. If you were to change this to `DEBUG`, it would produce much more detailed logging output.

As another example, let's say you wanted to set the output from the container-managed persistence engine to `DEBUG` level and to redirect it to a separate file, `cmp.log`, in order to analyze the generated SQL commands. You would add the following code to the `log4j.xml` file:

```
<appender name="CMP" class="org.jboss.logging.appender.RollingFileAppender">
  <errorHandler class="org.jboss.logging.util.OnlyOnceErrorHandler"/>
  <param name="File" value="{jboss.server.home.dir}/log/cmp.log"/>
  <param name="Append" value="false"/>
  <param name="MaxFileSize" value="500KB"/>
  <param name="MaxBackupIndex" value="1"/>

  <layout class="org.apache.log4j.PatternLayout">
    <param name="ConversionPattern" value="%d %-5p [%c] %m%n"/>
  </layout>
</appender>

<category name="org.jboss.ejb.plugins.cmp">
  <priority value="DEBUG" />
  <appender-ref ref="CMP"/>
</category>
```

This creates a new file appender and specifies that it should be used by the logger (or category) for the package `org.jboss.ejb.plugins.cmp`.

The file appender is set up to produce a new log file every day rather than producing a new one every time you restart the server or writing to a single file indefinitely. The current log file is `cmp.log`. Older files have the date they were written added to the name. You will notice that the `log` directory also contains HTTP request logs which are produced by the web container.

6.3. Security Service

The security domain information is stored in the file `login-config.xml` as a list of named security domains, each of which specifies a number of JAAS¹ login modules which are used for authentication purposes in that domain. When you want to use security in an application, you specify the name of the domain you want to use in the application's JBoss-specific deployment descriptors, `jboss.xml` and/or `jboss-web.xml`. We will quickly look at how to do this to secure the JMX Console application.

Almost every aspect of the JBoss server can be controlled through the JMX Console, so it is important to make sure that, at the very least, the application is password protected. Otherwise, any remote user could completely control your server. To protect it, we will add a security domain to cover the application.² This can be done in the `jboss-web.xml` file for the JMX Console, which can be found in `deploy/jmx-console.war/WEB-INF/` directory. Uncomment

¹ The Java Authentication and Authorization Service. JBoss uses JAAS to provide pluggable authentication modules. You can use the ones that are provided or write your own if have more specific requirements.

² If you installed JBoss using the Graphical Installer and set the JMX Security up, then you will not have to uncomment the sections, because they are already uncommented. Additionally, the admin password will be set up to whatever you had specified.

the `security-domain` in that file, as shown below.

```
<jboss-web>
  <security-domain>java:/jaas/jmx-console</security-domain>
</jboss-web>
```

This links the security domain to the web application, but it doesn't tell the web application what security policy to enforce, what URLs are we trying to protect, and who is allowed to access them. To configure this, go to the `web.xml` file in the same directory and uncomment the `security-constraint` that is already there. This security constraint will require a valid user name and password for a user in the `JBossAdmin` group.

```
<!--
  A security constraint that restricts access to the HTML JMX console
  to users with the role JBossAdmin. Edit the roles to what you want and
  uncomment the WEB-INF/jboss-web.xml/security-domain element to enable
  secured access to the HTML JMX console.
-->
<security-constraint>
  <web-resource-collection>
    <web-resource-name>HtmlAdaptor</web-resource-name>
    <description>
      An example security config that only allows users with the
      role JBossAdmin to access the HTML JMX console web application
    </description>
    <url-pattern>/*</url-pattern>
    <http-method>GET</http-method>
    <http-method>POST</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>JBossAdmin</role-name>
  </auth-constraint>
</security-constraint>
```

That's great, but where do the user names and passwords come from? They come from the `jmx-console` security domain we linked the application to. We have provided the configuration for this in the `conf/login-config.xml`.

```
<application-policy name="jmx-console">
  <authentication>
    <login-module
      code="org.jboss.security.auth.spi.UsersRolesLoginModule"
      flag="required">
      <module-option name="usersProperties">
        props/jmx-console-users.properties
      </module-option>
      <module-option name="rolesProperties">
        props/jmx-console-roles.properties
      </module-option>
    </login-module>
  </authentication>
```

```
</application-policy>
```

This configuration uses a simple file based security policy. The configuration files are found in the `conf/props` directory of your server configuration. The usernames and passwords are stored in the `conf/props/jmx-console-users.properties` file and take the form "username=password". To assign a user to the `JBossAdmin` group add "username=JBossAdmin" to the `jmx-console-roles.properties` file. The existing file creates an `admin` user with the password `admin`. For security, please either remove the user or change the password to a stronger one.

JBoss will re-deploy the JMX Console whenever you update its `web.xml`. You can check the server console to verify that JBoss has seen your changes. If you have configured everything correctly and re-deployed the application, the next time you try to access the JMX Console, it will ask you for a name and password.³

The JMX Console isn't the only web based management interface to JBoss. There is also the Web Console. Although it's a Java applet, the corresponding web application can be secured in the same way as the JMX Console. The Web Console is in the file `deploy/management/web-console.war`. The only difference is that the Web Console is provided as a simple WAR file instead of using the exploded directory structure that the JMX Console did. The only real difference between the two is that editing the files inside the WAR file is a bit more cumbersome.

6.4. Additional Services

The non-core, hot-deployable services are added to the `deploy` directory. They can be either XML descriptor files, `*-service.xml`, or JBoss Service Archive (SAR) files. SARs contain both the XML descriptor and additional resources the service requires (e.g. classes, library JAR files or other archives), all packaged up into a single archive.

Detailed information on all these services can be found in the *JBoss Enterprise Application Platform: Server Configuration Guide*, which also provides comprehensive information on server internals and the implementation of services such as JTA and the J2EE Connector Architecture (JCA).

7. The Web Container - Tomcat

JBoss Enterprise Application Platform comes with Tomcat as the default web container. The embedded Tomcat service is found in the `deploy/jboss-web.deployer` directory. All the necessary jar files needed by Tomcat can be found in there, as well as a `web.xml` file which provides a default configuration set for web applications.

³ Since the username and password are session variables in the web browser you may need to shut down your browser and come back in to see the login dialog come back up.

If you are already familiar with configuring Tomcat, have a look at the `server.xml`, which contains a subset of the standard Tomcat format configuration information. As it stands, this includes setting up the HTTP connector on the default port 8080, an AJP connector on port 8009 (can be used if you want to connect via a web server such as Apache) and an example of how to configure an SSL connector (commented out by default).

You shouldn't need to modify any of this other than for advanced use. If you've used Tomcat before as a stand-alone server you should be aware that things are a bit different when using the embedded service. JBoss is in charge and you shouldn't need to access the Tomcat directory at all. Web applications are deployed by putting them in the JBoss `deploy` directory and logging output from Tomcat can be found in the JBoss `log` directory.

EJB3 Caveats in JBoss Enterprise Application Platform 4.2.0.CP01

There are a number of implementation features that you should be aware of when developing applications for JBoss Enterprise Application Platform 4.2.0.CP01.

1. Unimplemented features

The Release Notes for JBoss Enterprise Application Platform 4.2.0.CP01 contain information on EJB3 features that are not yet implemented, or partially implemented. The Release Notes include links to issues in JIRA for information on workarounds and further details.

2. Referencing EJB3 Session Beans from non-EJB3 Beans

JBoss Enterprise Application Platform 5 will fully support the entire Java 5 Enterprise Edition specification. In the meantime JBoss Enterprise Application Platform 4.2.0.CP01 implements EJB3 functionality by way of an EJB MBean container running as a plugin in the JBoss Application Server. This has certain implications for application development.

The EJB3 plugin injects references to an EntityManager and @EJB references from one EJB object to another. However this support is limited to the EJB3 MBean and the JAR files it manages. Any JAR files which are loaded from a WAR (such as Servlets, JSF backing beans, and so forth) do not undergo this processing. The Java 5 Enterprise Edition standard specifies that a Servlet can reference a Session Bean through an @EJB annotated reference, however this is not implemented in JBoss Enterprise Application Platform 4.2.0.CP01.

In order to access an EJB3 Session Bean from a Servlet or JSF Backing Bean you will need to do one of two things:

1. **Without Seam - JNDI Lookup** . Without utilizing the Seam framework that is part of JBoss Enterprise Application Platform you will need to use an explicit JNDI lookup to access the EJB3 Session Bean. You can see an example of this being done in the `TodoBean.java` file in the `jsfejb3` example application, described in [Chapter 5, Sample JSF-EJB3 Application](#).

```
private TodoDaoInt getDao () {
    try {
        InitialContext ctx = new InitialContext();
        return (TodoDaoInt) ctx.lookup("jsfejb3/TodoDao/local");
    } catch (Exception e) {
        e.printStackTrace();
        throw new RuntimeException("couldn't lookup Dao", e);
    }
}
```

`ctx.lookup("jsfejb3/ToDoDao/local");` is the method used to reference the EJB3 Session Bean. The form is: *AppName/SessionBeanName/local*.

2. **With Seam - Leave it to the Seam Framework** . When you are using the Seam Framework you don't need to worry about this. Because the Seam framework manages the interaction of Beans anyway, it already automates this type of interaction.

Refer to [Chapter 6, Using Seam](#) for a more detailed explanation of achieving this using the Seam framework.

About the Example Applications

In this guide, we make use of a simple web application to illustrate the use of JSF-EJB3 components. We then illustrate how to use Seam to integrate the JSF and EJB3 components. The example applications (source code) come with this guide and you can find them located in the `JBOSS_DIST/doc/examples` directory. You can also download the sample applications from <http://www.redhat.com/docs/manuals/jboss>. We use two examples in this book:

- A simple "TODO" application to create, view and edit tasks - implemented using JSF and EJB3;
- The same application using the SEAM framework.

If you installed the documentation on your hard drive, then the first example can be found in the `JBOSS_DIST/doc/examples/jsfejb3` directory. We will see how to build this example using the `build.xml` file present here and also how to deploy the application. We will also cover in detail the working of the `.java`, `.xml` and `.properties` files.

The second example used in this guide can be found in the `JBOSS_DIST/doc/examples/seamejb3` directory. Using a simple "TODO" application we will illustrate how Seam ties together the database, the web interface and the EJB3 business logic in a web application. We will use the `build.xml` file present here to compile and build our Seam application.

Within the `JBOSS_DIST/doc/examples/<seamejb3 | jsfejb3>` directory, you will find the following sub-directories:

- **src**: contains the Java source code files.
- **view**: contains the web pages.
- **resources**: contains all the configuration files used.

1. Install Ant

To compile and package the examples, you must have Apache Ant 1.6+ installed in your machine. You can download it from <http://ant.apache.org> and have it installed in few steps:

- Unzip the downloaded file to the directory of your choice.
- Create an environment variable called `ANT_HOME` pointing to the Ant installation directory. You can do this by adding the following line to your `.bashrc` file (substituting with the actual location of the ant directory on your system):

```
export ANT_HOME=/home/user/apache-ant-1.7.0
```

On Windows you do this by opening the Control Panel from the Start Menu, switching it to classic view if necessary, then opening System/Advanced/Environment Variables. Create a new variable, call it `ANT_HOME` and set it to be the ant directory.

- Add `$ANT_HOME/bin` to the system path to be able to run `ant` from the command line. You can do this by adding the following line to your `.bashrc` file:

```
export PATH=$PATH:$ANT_HOME/bin
```

On Windows you do this by opening the Control Panel from the Start Menu, switching it to classic view if necessary, then editing the `PATH` environment variable found in System/Advanced/Environment Variables/System Variables/Path. Add a semicolon and the path to the ant `bin` directory.

- Verify your Ant installation. To do this type `ant -version` at the command prompt. Your output should look something like this:

```
Apache Ant version 1.7.0 compiled on December 13 2006
```

Sample JSF-EJB3 Application

We use a simple "TODO" application to show how JSF and EJB3 work together in a web application. The "TODO" application works like this: You can create a new 'todo' task item using the "Create" web form. Each 'todo' item has a 'title' and a 'description'. When you submit the form, the application saves your task to a relational database. Using the application, you can view all 'todo' items, edit/delete an existing 'todo' item and update the task in the database.

The sample application comprises the following components:

- Entity objects - These objects represent the data model; the properties in the object are mapped to column values in relational database tables.
- JSF web pages - The web interface used to capture input data and display result data. The data fields on these web pages are mapped to the data model via the JSF Expression Language (EL).
- EJB3 Session Bean - This is where the functionality is implemented. We make use of a Stateless Session Bean.

1. Data Model

Lets take a look at the contents of the Data Model represented by the `Todo` class in the `Todo.java` file. Each instance of the `Todo` class corresponds to a row in the relational database table. The 'Todo' class has three properties: `id`, `title` and `description`. Each of these correspond to a column in the database table.

The 'Entity class' to 'Database Table' mapping information is specified using EJB3 Annotations in the 'Todo' class. This eliminates the need for XML configuration and makes it a lot clearer. The `@Entity` annotation defines the `Todo` class as an Entity Bean. The `@Id` and `@GeneratedValue` annotations on the `id` property indicate that the `id` column is the primary key and that the server automatically generates its value for each `Todo` object saved into the database.

```
@Entity
public class Todo implements Serializable {

    private long id;
    private String title;
    private String description;

    public Todo () {
        title = "";
        description = "";
    }

    @Id @GeneratedValue
```

```
public long getId() { return id;}
public void setId(long id) { this.id = id; }

public String getTitle() { return title; }
public void setTitle(String title) {this.title = title;}

public String getDescription() { return description; }
public void setDescription(String description) {
    this.description = description;
}
}
```

2. JSF Web Pages

In this section we will show you how the web interface is defined using JSF pages. We will also see how the data model is mapped to the web form using JSF EL. Using the `#{...}` notation to reference Java objects is called **JSF EL** (JSF Expression Language). Lets take a look at the pages used in our application:

- **index.xhtml**: This page displays two options: 1. Create New Todo 2. Show all Todos. When you click on the Submit button the corresponding action is invoked.

```
<h:form>
<ul>
  <li><h:commandLink type="submit" value="Create New Todo"
action="create"/></li>
  <li><h:commandLink type="submit" value="Show All Todos"
action="todos"/></li>
</ul>
</h:form>
```

- **create.xhtml**: When you try to create a new task, this JSF page captures the input data. We use the `todoBean` to back the form input text fields. The `#{todoBean.todo.title}` symbol refers to the "title" property of the "todo" object in the "TodoBean" class. The `#{todoBean.todo.description}` symbol refers to the "description" property of the "todo" object in the "TodoBean" class. The `#{todoBean.persist}` symbol refers to the "persist" method in the "TodoBean" class. This method creates the "Todo" instance with the input data (title and description) and persists the data.

```
<h:form id="create">
<table>
  <tr>
    <td>Title:</td>
    <td>
```

```

        <h:inputText id="title" value="#{todoBean.todo.title}" size="15">
            <f:validateLength minimum="2"/>
        </h:inputText>
    </td>
</tr>
<tr>
    <td>Description:</td>
    <td>
        <h:inputTextarea id="description"
value="#{todoBean.todo.description}">
            <f:validateLength minimum="2" maximum="250" />
        </h:inputTextarea>
    </td>
</tr>
</table>
<h:commandButton type="submit" id="create" value="Create"
    action="#{todoBean.persist}" />
</h:form>

```

Figure 5.1, "The "Create Todo" web page" shows the "Create Todo" web page with the input fields mapped to the data model.

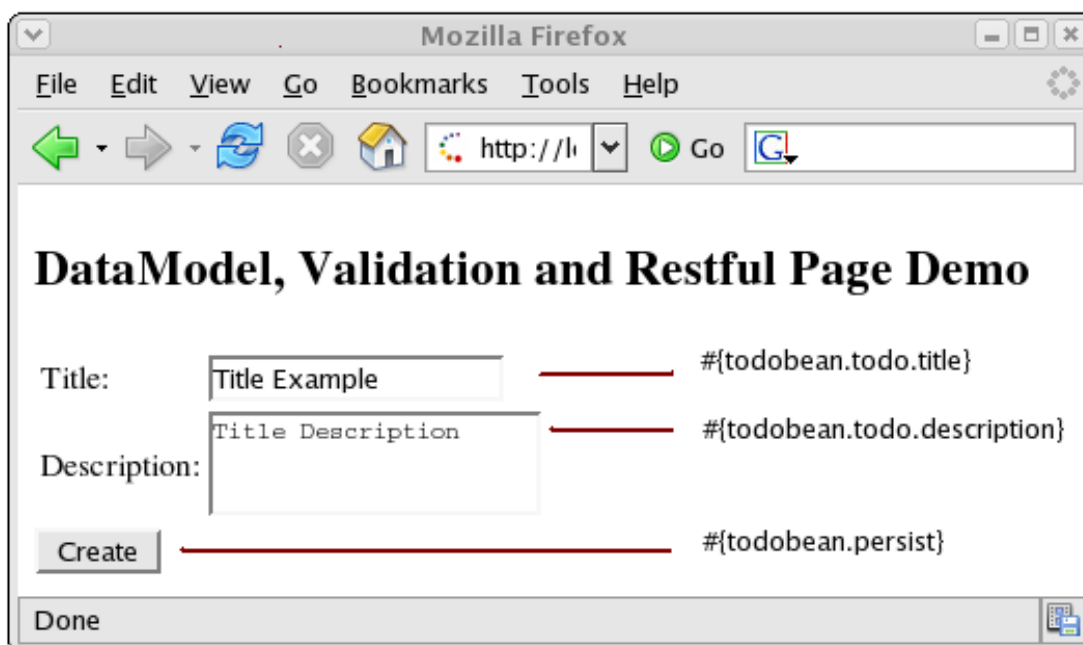


Figure 5.1. The "Create Todo" web page

- **todos.xhtml:** This page displays the list of all "todos" created. There is also an option to choose a "todo" item for 'edit' or 'delete'.

The list of all 'todos' is fetched by `#{todoBean.todos}` symbol referring to the 'getTodos()' property in the 'TodoBean' class. The JSF `dataTable` iterates through the list and displays

each `Todo` object in a row. The 'Edit' option is available across each row. The `#{todo.id}` symbol represents the "id" property of the "todo" object.

```
<h:form>
<h:dataTable value="#{todoBean.todos}" var="todo">
  <h:column>
    <f:facet name="header">Title</f:facet>
    #{todo.title}
  </h:column>
  <h:column>
    <f:facet name="header">Description</f:facet>
    #{todo.description}
  </h:column>
  <h:column>
    <a href="edit.faces?tid=#{todo.id}">Edit</a>
  </h:column>
</h:dataTable>
<center>
  <h:commandButton action="create"
    value="Create New Todo" type="submit"/>
</center>
</h:form>
```

Figure 5.2, "The "Show All Todos" web page " shows the "Show All Todos" web page with the data fields mapped to the data model.



Figure 5.2. The "Show All Todos" web page

- **edit.xhtml**: This page allows you to edit the "todo" item's 'title' and 'description' properties. The `#{todoBean.update}` and `#{todoBean.delete}` symbols represent the "update" and "delete" methods in the "TodoBean" class.

```
<h2>Edit #{todoBean.todo.title}</h2>
<h:form id="edit">
<input type="hidden" name="tid" value="#{todoBean.todo.id}"/>
<table>
  <tr>
    <td>Title:</td>
    <td>
      <h:inputText id="title" value="#{todoBean.todo.title}" size="15">
        <f:validateLength minimum="2"/>
      </h:inputText>
    </td>
  </tr>
  <tr>
    <td>Description:</td>
    <td>
      <h:inputTextarea id="description"
value="#{todoBean.todo.description}">
        <f:validateLength minimum="2" maximum="250"/>
      </h:inputTextarea>
    </td>
  </tr>
</table>
<h:commandButton type="submit" id="update" value="Update"
  action="#{todoBean.update}"/>
<h:commandButton type="submit" id="delete" value="Delete"
  action="#{todoBean.delete}"/>
</h:form>
```

Figure 5.3, "The "Edit Todo" web page " shows the "Edit Todo" web page with the mapping to the data model.

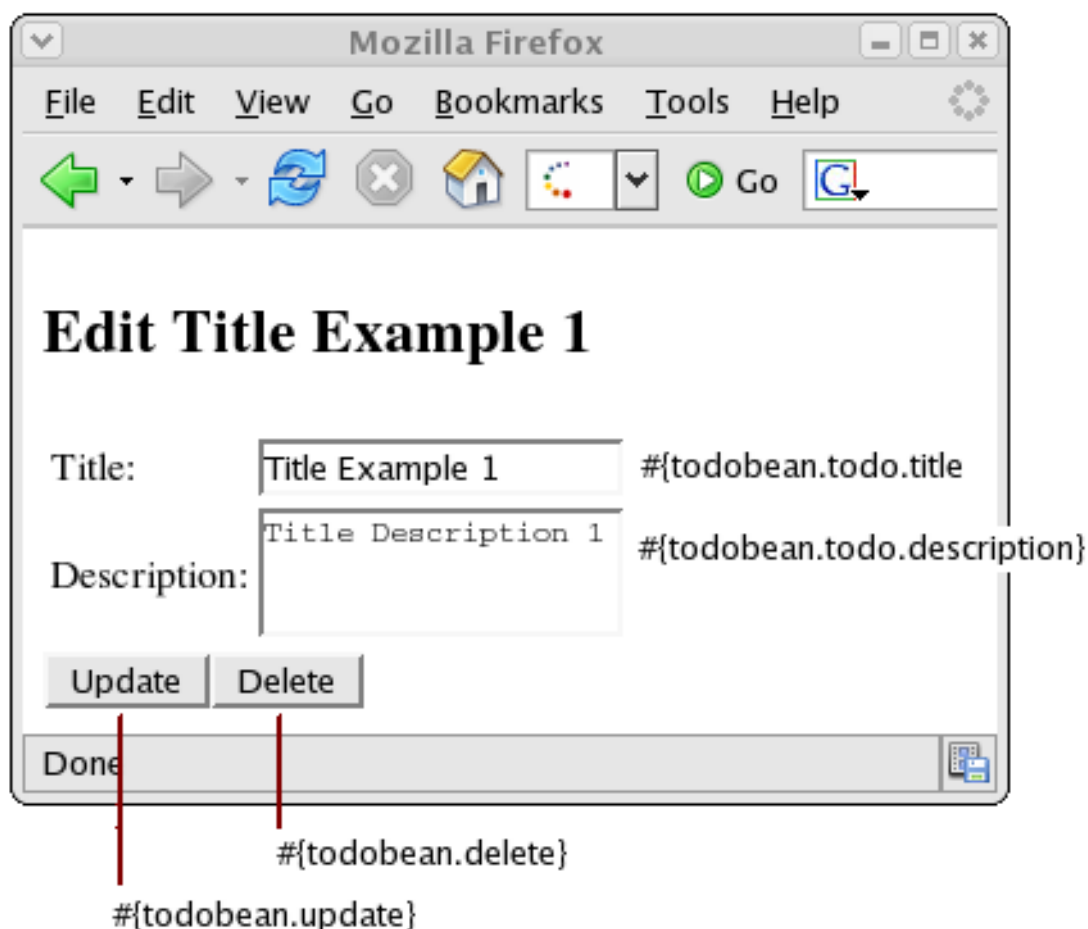


Figure 5.3. The "Edit Todo" web page



Note

We have used XHTML pages in the sample applications because we recommend using Facelets instead of JSP to render JSF view pages.

3. EJB3 Session Beans

EJB 3.0 is one of the major improvements introduced with Java EE 5.0. It aims at reducing the complexity of older versions of EJB and simplifies Enterprise Java development and deployment. You will notice that to declare a class as a 'Session Bean' you simply have to annotate it. Using annotations eliminates the complexity involved with too many deployment descriptors. Also the only interface an EJB3 Session Bean requires is a business interface that declares all the business methods that must be implemented by the bean.

We will explore the two important source files associated with the Bean implementation in our application: `TodoDaoInt.java` and `TodoDao.java`.

- **Business interface:** `TodoDaoInt.java`

We define here the methods that need to be implemented by the bean implementation class. Basically, the business methods that will be used in our application are defined here.

```
public interface TodoDaoInt {

    public void persist (Todo todo);
    public void delete (Todo todo);
    public void update (Todo todo);

    public List <Todo> findTodos ();
    public Todo findTodo (String id);
}
```

- **Stateless Session Bean:** `TodoDao.java`

The `@Stateless` annotation marks the bean as a stateless session bean. In this class, we need to access the Entity bean `Todo` defined earlier. For this we need an `EntityManager`. The `@PersistenceContext` annotation tells the JBoss Server to inject an entity manager during deployment.

```
@Stateless
public class TodoDao implements TodoDaoInt {

    @PersistenceContext
    private EntityManager em;

    public void persist (Todo todo) {
        em.persist (todo);
    }

    public void delete (Todo todo) {
        Todo t = em.merge (todo);
        em.remove( t );
    }

    public void update (Todo todo) {
        em.merge (todo);
    }

    public List <Todo> findTodos () {
        return (List <Todo>) em.createQuery("select t from Todo t")
            .getResultList();
    }

    public Todo findTodo (String id) {
```

```
        return (Todo) em.find(Todo.class, Long.parseLong(id));
    }
}
```

4. Configuration and Packaging

We will build the sample application using Ant and explore the configuration and packaging details. If you haven't installed Ant yet, do so now.

4.1. Building The Application

Let's look at building the example application and then explore the configuration files in detail.

In [Chapter 4, About the Example Applications](#), we looked at the directory structure of the `jsfejb3` sample application. At the command line, go to the `jsfejb3` directory. There you will see a `build.xml` file. This is our Ant build script for compiling and packaging the archives. To build the application, you need to first of all edit the `build.xml` file and edit the value of `jboss-dist` to reflect the location where the JBoss Application Server is installed. Once you have done this, just type the command `ant` and your output should look like this:

```
[vrenish@vinux jsfejb3]$ ant
Buildfile: build.xml

compile:
  [mkdir] Created dir:
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3/build/classes
  [javac] Compiling 4 source files to
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3
/build/classes
  [javac] Note:
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3/src/ToDoDao.java uses
unchecked or unsafe operations.
  [javac] Note: Recompile with -Xlint:unchecked for details.

war:
  [mkdir] Created dir:
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3/build/jars
  [war] Building war:
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3/build/jars/
app.war

ejb3jar:
  [jar] Building jar:
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3/build/jars/
app.jar

ear:
  [ear] Building ear:
/home/vrenish/jboss-eap-4.2/doc/examples/jsfejb3/build/jars/
```

```
jsfejb3.ear

main:

BUILD SUCCESSFUL
Total time: 2 seconds
(vrenish@vinux jsfejb3)$
```

If you get the BUILD SUCCESSFUL message, you will find a newly created `build` directory with 2 sub-directories in it:

- **classes**: containing the compiled class files.
- **jars**: containing three archives - `app.jar`, `app.war` and `jsfejb3.ear`.
 - `app.jar` : EJB code and descriptors.
 - `app.war` : web application which provides the front end to allow users to interact with the business components (the EJBs). The web source (HTML, images etc.) contained in the `jsfejb3/view` directory is added unmodified to this archive. The Ant task also adds the `WEB-INF` directory that contains the files which aren't meant to be directly accessed by a web browser but are still part of the web application. These include the deployment descriptors (`web.xml`) and extra jars required by the web application.
 - `jsfejb3.ear` : The EAR file is the complete application, containing the EJB modules and the web module. It also contains an additional descriptor, `application.xml`. It is also possible to deploy EJBs and web application modules individually but the EAR provides a convenient single unit.

4.2. Configuration Files

Now that we have built the application, lets take a closer look at some of the important Configuration files. We have built the final archive ready for deployment - `jsfejb3.ear`. The contents of your EAR file should look like this:

```
jsfejb3.ear
|+ app.jar    // contains the EJB code
  |+ import.sql
  |+ Todo.class
  |+ TodoDao.class
  |+ TodoDaoInt.class
  |+ META-INF
    |+ persistence.xml
|+ app.war    // contains web UI
  |+ index.html
  |+ index.xhtml
  |+ create.xhtml
  |+ edit.xhtml
```

```
|+ todos.xhtml
|+ TodoBean.class
|+ style.css
|+ META-INF
|+ WEB-INF
  |+ faces-config.xml
  |+ navigation.xml
  |+ web.xml
|+ META-INF // contains the descriptors
  |+ application.xml
  |+ jboss-app.xml
```

- `application.xml`: This file lists the JAR files in the EAR (in our case `app.jar`) and tells the JBoss server what files to look for and where. The root URL for the application is also specified in this file as 'context-root'.

```
<application>
  <display-name>Sample Todo</display-name>
  <module>
    <web>
      <web-uri>app.war</web-uri>
      <context-root>/jsfejb3</context-root>
    </web>
  </module>
  <module>
    <ejb>app.jar</ejb>
  </module>
</application>
```

- `jboss-app.xml`: Every EAR application should specify a unique string name for the class loader. In our case, we use the application name 'jsfejb3' as the class loader name.

```
<jboss-app>
  <loader-repository>
    jsfejb3:archive=jsfejb3.ear
  </loader-repository>
</jboss-app>
```

- `app.jar`: This contains EJB3 Session Bean and Entity Bean classes and the related configuration files. In addition, the `persistence.xml` file configures the back-end data source (in our case the default HSQL database) for the `EntityManager`.

```
<persistence>
  <persistence-unit name="helloworld">
    <provider>org.hibernate.ejb.HibernatePersistence</provider>
    <jta-data-source>java:/DefaultDS</jta-data-source>
```

```

    <properties>
      <property name="hibernate.dialect"
value="org.hibernate.dialect.HSQLDialect"/>
      <property name="hibernate.hbm2ddl.auto" value="create-drop"/>
    </properties>
  </persistence-unit>
</persistence>

```

- `app.war`: This contains the Web UI files packaged according to the Web Application Archive (WAR) specification. It contains all the web pages and the required configuration files. The `web.xml` file is an important file for all JAVA EE web applications. It is the web deployment descriptor file. The `faces-config.xml` file is the configuration file for JSF. The `navigation.xml` file contains the rules for JSF page navigation.

```

//faces-config.xml
<faces-config>
  <application>
    <view-handler>
      com.sun.facelets.FaceletViewHandler
    </view-handler>
  </application>
  <managed-bean>
    <description>Dao</description>
    <managed-bean-name>todoBean</managed-bean-name>
    <managed-bean-class>TodoBean</managed-bean-class>
    <managed-bean-scope>request</managed-bean-scope>
  </managed-bean>
</faces-config>

```

5. The Database

5.1. Creating the Database Schema

To pre-populate the database, we have supplied SQL Code (`import.sql`) to run with HSQL in the `examples/jsfejb3/resources` directory. When you build the application using Ant, this is packaged in the `app.jar` file within the `jsfejb3.ear` file. When the application is deployed, you should be able to view the pre-populated data.

5.2. The HSQL Database Manager Tool

Just as a quick aside at this point, start up the JMX console application and click on the `service=Hypersonic` link which you'll find under the section `jboss`. If you can't find this, make sure the Hypersonic service is enabled in the `hsqldb-ds.xml` file.

This will take you to the information for the Hypersonic service MBean. Scroll down to the bottom of the page and click the `invoke` button for the `startDatabaseManager()` operation.

This starts up the HSQL Manager, a Java GUI application which you can use to manipulate the database directly.

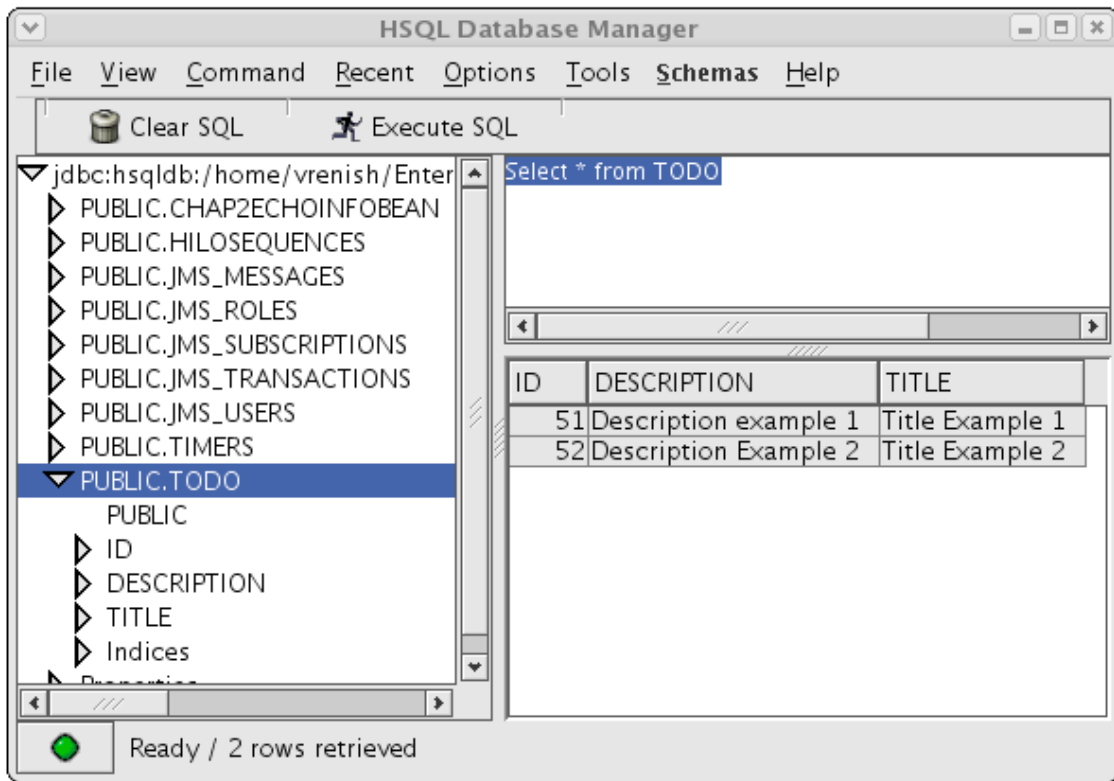


Figure 5.4. The HSQL Database Manger

6. Deploying the Application

Deploying an application in JBoss is simple and easy. You just have to copy the EAR file to the `deploy` directory in the 'server configuration' directory of your choice. Here, we will deploy it to the 'default' configuration, so we copy the EAR file to the `JBOSS_DIST/jboss-as/server/default/deploy` directory.

You should see something close to the following output from the server:

```
15:32:23,997 INFO [EARDeployer] Init J2EE application:
file:/home/vrenish/jboss-eap-4.2
/jboss-as/server/default/deploy/jsfejb3.ear
15:32:24,212 INFO [JmxKernelAbstraction] creating wrapper delegate for:
org.jboss.ejb3.
entity.PersistenceUnitDeployment
15:32:24,213 INFO [JmxKernelAbstraction] installing MBean:
persistence.units:ear=
jsfejb3.ear,jar=app.jar,unitName=helloworld with dependencies:
15:32:24,213 INFO [JmxKernelAbstraction]
jboss.jca:name=DefaultDS,service=
```

```
DataSourceBinding
15:32:24,275 INFO [PersistenceUnitDeployment] Starting persistence unit
persistence.
units:ear=jsfejb3.ear,jar=app.jar,unitName=helloworld
15:32:24,392 INFO [Ejb3Configuration] found EJB3 Entity bean: Todo
15:32:24,450 WARN [Ejb3Configuration] Persistence provider caller does not
implements
the EJB3 spec correctly. PersistenceUnitInfo.getNewTempClassLoader() is
null.
15:32:24,512 INFO [Configuration] Reading mappings from resource :
META-INF/orm.xml
15:32:24,512 INFO [Ejb3Configuration] [PersistenceUnit: helloworld] no
META-INF/orm.xml
found
15:32:24,585 INFO [AnnotationBinder] Binding entity from annotated class:
Todo
15:32:24,586 INFO [EntityBinder] Bind entity Todo on table Todo
.
.
.
.
15:32:26,311 INFO [SchemaExport] Running hbm2ddl schema export
15:32:26,312 INFO [SchemaExport] exporting generated schema to database
15:32:26,314 INFO [SchemaExport] Executing import script: /import.sql
15:32:26,418 INFO [SchemaExport] schema export complete
15:32:26,454 INFO [NamingHelper] JNDI InitialContext
properties:{java.naming.factory.
initial=org.jnp.interfaces.NamingContextFactory,
java.naming.factory.url.pkgs=org.jboss.
naming:org.jnp.interfaces}
15:32:26,484 INFO [JmxKernelAbstraction] creating wrapper delegate for:
org.jboss.ejb3.
stateless.StatelessContainer
15:32:26,485 INFO [JmxKernelAbstraction] installing MBean:
jboss.j2ee:ear=jsfejb3.ear,
jar=app.jar,name=TodoDao,service=EJB3 with dependencies:
15:32:26,513 INFO [JmxKernelAbstraction]
persistence.units:ear=jsfejb3.ear,
jar=app.jar,unitName=helloworld
15:32:26,557 INFO [EJBContainer] STARTED EJB: TodoDao ejbName: TodoDao
15:32:26,596 INFO [EJB3Deployer] Deployed:
file:/home/vrenish/jboss-eap-4.2/jboss-as/
server/default/tmp/deploy/
tmp33761jsfejb3.ear-contents/app.jar
15:32:26,625 INFO [TomcatDeployer] deploy, ctxPath=/jsfejb3,
warUrl=../tmp/deploy/
tmp33761jsfejb3.ear-contents/app-exp.war/
15:32:26,914 INFO [EARDeployer] Started J2EE application:
file:/home/vrenish/jboss-eap-
4.2/jboss-as/server/default/deploy/jsfejb3.ear
```

If there are any errors or exceptions, make a note of the error message. Check that the EAR is complete and inspect the WAR file and the EJB jar files to make sure they contain all the necessary components (classes, descriptors etc.).

You can safely redeploy the application if it is already deployed. To undeploy it you just have to remove the archive from the `deploy` directory. There's no need to restart the server in either case. If everything seems to have gone OK, then point your browser at the application URL.

<http://localhost:8080/jsfejb3>

You will be forwarded to the application main page. *Figure 5.5, "Sample TODO"* shows the sample application in action.



Figure 5.5. Sample TODO

Using Seam

JBoss Seam is a framework that provides the glue between the new EJB3 and JSF frameworks that are part of the Java EE 5.0 standard. In fact, the name Seam refers to the seamless manner in which it enables developers to use these two frameworks in an integrated manner. Seam automates many of the common tasks, and makes extensive use of annotations to reduce the amount of xml code that needs to be written. The overall effect is to significantly reduce the total amount of coding that needs to be done.

We have included two versions of the example application, one coded using EJB3 / JSF without using Seam, and one using Seam, to demonstrate clearly the difference in application development using the Seam framework.

1. Data Model

In the previous chapter we looked at the Data Model used in the EJB3/JSF implementation of this sample application. Let's start off our examination of the Seam implementation in the same way, by examining how the Data Model is implemented. This is done in the `Todo.java` file.

```
@Entity
@Name("todo")
public class Todo implements Serializable {

    private long id;
    private String title;
    private String description;

    public Todo () {
        title = "";
        description = "";
    }

    @Id @GeneratedValue
    public long getId() { return id;}
    public void setId(long id) { this.id = id; }

    @NotNull
    public String getTitle() { return title; }
    public void setTitle(String title) {this.title = title;}

    @NotNull
    @Length(max=250)
    public String getDescription() { return description; }
    public void setDescription(String description) {
        this.description = description;
    }
}
```

The `@Entity` annotation defines the class as an EJB3 session bean, and tells the container to

map the `Todo` class to a relational database table. Each property of the class will become a column in the table. Each instance of the class will become a row in this table. Since we have not used the `@Table` annotation, Seam's "configuration by exception" default will name the table after the class.

`@Entity` and `@Table` are both EJB3 annotations, and are not specific to Seam. It is possible to use Seam completely with POJOs (Plain Old Java Objects) without any EJB3-specific annotations. However, EJB3 brings a lot of advantages to the table, including container managed security, message-driven components, transaction and component level persistence context, and `@PersistenceContext` injection, which we will encounter a little further on.

The `@Name` annotation is specific to Seam, and defines the string name for Seam to use to register the Entity Bean. This will be the default name for the relational database table. Each component in a Seam application must have a unique name. In the other components in the Seam framework, such as JSF web pages and session beans, you can reference the managed `Todo` bean using this name. If no instance of this class exists when it is referenced from another component, then Seam will instantiate one.

The `@Id` annotation defines a primary key `id` field for the component. `@GeneratedValue` specifies that the server will automatically generate this value for the component when it is saved to the database.

Seam provides support for model-based constraints defined using Hibernate Validator, although Hibernate does not have to be the object persister used. The `@NotNull` annotation is a validation constraint that requires this property to have a value before the component can be persisted into the database. Using this annotation allows the validation to be enforced by the JSF code at the view level, without having to specify the exact validation constraint in the JSF code.

At this point the only apparent difference between the Seam version and the EJB3/JSF version of the app is the inclusion of the validator annotation `@NotNull`, and the `@Name` annotation. However, while the EJB3/JSF version of this application requires a further `TodoBean` class to be manually coded and managed in order to handle the interaction between the `Todo` class and the web interface, when using Seam the Seam framework takes care of this work for us. We'll see how this is done in practice as we examine the implementation of the user interface.

2. JSF Web Pages - `index.xhtml` and `create.xhtml`

The `index.xhtml` file used is the same as in the EJB3/JSF example.

`create.xhtml` begins to reveal the difference that coding using the Seam framework makes.

```
<h:form id="create">

  <f:facet name="beforeInvalidField">
    <h:graphicImage styleClass="errorImg" value="error.png" />
  </f:facet>
  <f:facet name="afterInvalidField">
    <s:message styleClass="errorMsg" />
  </f:facet>
```

```
<f:facet name="aroundInvalidField">
  <s:div styleClass="error" />
</f:facet>

<s:validateAll>

<table>

  <tr>
    <td>Title:</td>
    <td>
      <s:decorate>
        <h:inputText id="title" value="#{todo.title}" size="15" />
      </s:decorate>
    </td>
  </tr>

  <tr>
    <td>Description:</td>
    <td>
      <s:decorate>
        <h:inputTextarea id="description" value="#{todo.description}" />
      </s:decorate>
    </td>
  </tr>

</table>

</s:validateAll>

<h:commandButton type="submit" id="create" value="Create"
  action="#{todoDao.persist}" />
</h:form>
```

The first thing that is different here is the Java Server Facelet code at the beginning, which works with the `@NotNull` validation constraint of our `todo` class to enforce and indicate invalid input to the user.

Also notice here that rather than requiring the use of a `TodoBean` class as we did in the EJB3/JSF example we back the form directly with a `Todo` entity bean. When this page is called, JSF asks Seam to resolve the variable `todo` due to JSF EL references such as `#{todo.title}`. Since there is no value already bound to that variable name, Seam will instantiate an entity bean of the `todo` class and return it to JSF, after storing it in the Seam context. The Seam context replaces the need for an intermediary bean.

The form input values are validated against the Hibernate Validator constraints specified in the `todo` class. JSF will redisplay the page if the constraints are violated, or it will bind the form input values to the `Todo` entity bean.

Entity beans shouldn't do database access or transaction management, so we can't use the `Todo` entity bean as a JSF action listener. Instead, creation of a new `todo` item in the database is

accomplished by calling the `persist` method of a `TodoDao` session bean. When JSF requests Seam to resolve the variable `todoDao` through the JSF EL expression `#{todoDao.persist}`, Seam will either instantiate an object if one does not already exist, or else pass the existing stateful `todoDao` object from the Seam context. Seam will intercept the `persist` method call and inject the `todo` entity from the session context.

Let's have a look at the `TodoDao` class (defined in `TodoDao.java`) to see how this injection capability is implemented.

3. Data Access using a Session Bean

Let's go through a listing of the code for the `TodoDao` class.

```
@Stateful
@Name("todoDao")
public class TodoDao implements TodoDaoInt {

    @In (required=false) @Out (required=false)
    private Todo todo;

    @PersistenceContext (type=EXTENDED)
    private EntityManager em;

    // Injected from pages.xml
    Long id;

    public String persist () {
        em.persist (todo);
        return "persisted";
    }

    @DataModel
    private List <Todo> todos;

    @Factory("todos")
    public void findTodos () {
        todos = em.createQuery("select t from Todo t")
                .getResultList();
    }

    public void setId (Long id) {
        this.id = id;

        if (id != null) {
            todo = (Todo) em.find(Todo.class, id);
        } else {
            todo = new Todo ();
        }
    }

    public Long getId () {
        return id;
    }

    public String delete () {
```

```

        em.remove( todo );
        return "removed";
    }

    public String update () {
        return "updated";
    }

    @Remove @Destroy
    public void destroy() {}

}

```

First of all notice that this is a stateful session bean. Seam can use both stateful and stateless session beans, the two most common types of EJB3 beans.

The `@In` and `@Out` annotations define an attribute that is injected by Seam. The attribute is injected to this object or from this object to another via a Seam context variable named `todo`, a reference to the Seam registered name of our `Todo` class defined in `Todo.java`.

The `@PersistenceContext` annotation injects the EJB3 Entity manager, allowing this object to persist objects to the database. Because this is a stateful session bean and the `PersistenceContext` type is set to `EXTENDED`, the same Entity Manager instance is used until the Remove method of the session bean is called. The database to be used (a persistence-unit) is defined in the file `resources/META-INF/persistence.xml`

Note that this session bean has simultaneous access to context associated with web request (the form values of the `todo` object), and state held in transactional resources (the `EntityManager`). This is a break from traditional J2EE architectures, but Seam does not force you to work this way. You can use more traditional forms of application layering if you wish.

The `@DataModel` annotation initializes the `todos` property, which will be outjected or "exposed" to the view. The `@Factory` annotated method performs the work of generating the `todos` list, and is called by Seam if it attempts to access the exposed `DataModel` property and finds it to be null. Notice the absence of property access methods for the `todos` property. Seam takes care of this for you automatically.

Let's take a look at the JSF code that we use for displaying and editing the list of todos, to get an idea of how to use these interfaces in practice.

4. JSF Web Pages - todos.xhtml and edit.xhtml

Using the `DataModel` exposed property of the Session Bean it becomes trivial to produce a list of todos:

```

<h:form>

<h:dataTable value="#{todos}" var="todo">
  <h:column>
    <f:facet name="header">Title</f:facet>

```

```
        #{todo.title}
    </h:column>
    <h:column>
        <f:facet name="header">Description</f:facet>
        #{todo.description}
    </h:column>
    <h:column>
        <a href="edit.seam?tid=#{todo.id}">Edit</a>
    </h:column>
</h:dataTable>

<center>
    <h:commandButton action="create"
        value="Create New Todo" type="submit"/>
</center>

</h:form>
```

When the JSF variable resolver encounters `{#todos}` and requests `todos`, Seam finds that there is no "todos" component in the current scope, so it calls the `@Factory("todos")` method to make one. The `todos` object is then outjected once the factory method is done since it is annotated with the `@DataModel` annotation.

Constructing the view for the edit page is similarly straight forward:

```
<h:form id="edit">

    <f:facet name="beforeInvalidField">
        <h:graphicImage styleClass="errorImg" value="error.png"/>
    </f:facet>
    <f:facet name="afterInvalidField">
        <s:message styleClass="errorMsg" />
    </f:facet>
    <f:facet name="aroundInvalidField">
        <s:div styleClass="error"/>
    </f:facet>

    <s:validateAll>

    <table>

        <tr>
            <td>Title:</td>
            <td>
                <s:decorate>
                    <h:inputText id="title" value="{#todo.title}" size="15"/>
                </s:decorate>
            </td>
        </tr>

        <tr>
            <td>Description:</td>
            <td>
                <s:decorate>
```

```

        <h:inputTextarea id="description" value="#{todo.description}"/>
    </s:decorate>
</td>
</tr>

</table>

</s:validateAll>

<h:commandButton type="submit" id="update" value="Update"
    action="#{todoDao.update}"/>

<h:commandButton type="submit" id="delete" value="Delete"
    action="#{todoDao.delete}"/>

</h:form>

```

Here we see the same factors in play. JSF validation code taking advantage of the validation constraints defined in our Entity Bean, and the use of the `todoDao` Session Bean's `update` and `delete` methods to update the database.

The call from `todos.xhtml:edit.seam?tid=#{todo.id}` causes Seam to create a `todoDao` and set its `id` property to `tid`. Setting its `id` property causes the `todoDao` to retrieve the appropriate record from the database.

The functionality that allows the edit page to be called with a parameter in this way is implemented through `pages.xml`. Let's have a look at the `pages.xml` file and how it is used by Seam applications.

5. Xml Files

Seam drastically reduces the amount of xml coding that needs to be done. One file that is of interest is the `pages.xml`, packaged in the `app.war` file's `WEB-INF` directory. This file is available in the `resources/WEB-INF` directory in the source code bundle. The `pages.xml` file is used to define page descriptions including Seam page parameters (HTTP `GET` parameters), page actions, page navigation rules, error pages etc. Among other things it can be used in a Seam application to define exception handlers and redirections.

In the case of our sample application we are using it to define a Seam page parameter. The `pages.xml` in this example contains the following code:

```

<page view-id="/edit.xhtml">
    <param name="tid" value="#{todoDao.id}"
        converterId="javax.faces.Long"/>
</page>

```

This defines a parameter named `tid` for the `edit.xhtml` page. When the `edit.xhtml` page is loaded, the HTTP `GET` request parameter `tid` is converted to a `Long` value and assigned to the `id` property of the `todoDao` object. You can have as many page parameters as required to bind

HTTP `GET` request parameters to the back-end components in your application.

6. Further Information

This completes our walkthrough of the sample Seam application. For further, detailed information on developing applications using the Seam framework, please refer to the "Seam Reference Guide".

Using other Databases

In the previous chapters, we've just been using the JBoss default datasource in our applications. This is provided by the embedded HSQL database instance and is bound to the JNDI name `java:/DefaultDS`. Having a database included with JBoss is very convenient for running examples and HSQL is adequate for many purposes. However, at some stage you will want to use another database, either to replace the default datasource or to access multiple databases from within the server.

1. DataSource Configuration Files

DataSource configuration file names end with the suffix `-ds.xml` so that they will be recognized correctly by the JCA deployer. The `docs/example/jca` directory contains sample files for a wide selection of databases and it is a good idea to use one of these as a starting point. For a full description of the configuration format the best place to look is the DTD file `docs/dtd/jboss-ds_1_5.dtd`. Additional documentation on the files and the JBoss JCA implementation can also be found in the *JBoss 4 Application Server Guide*.

Local transaction datasources are configured using the `local-tx-datasource` element and XA-compliant ones using `xa-tx-datasource`. The example file `generic-ds.xml` shows how to use both types and also some of the other elements that are available for things like connection pool configuration. Examples of both local and XA configurations are available for Oracle, DB2 and Informix.

If you look at the example files `firebird-ds.xml`, `facets-ds.xml` and `sap3-ds.xml`, you'll notice that they have a completely different format, with the root element being `connection-factories` rather than `datasources`. These use an alternative, more generic JCA configuration syntax used with a pre-packaged JCA resource adapter. The syntax is not specific to datasource configuration and is used, for example, in the `jms-ds.xml` file to configure the JMS resource adapter.

Next, we'll work through some step-by-step examples to illustrate what's involved setting up a datasource for a specific database.

2. Using MySQL as the Default DataSource

MySQL is a one of the most popular open source databases around and is used by many prominent organizations from Yahoo to NASA. The official JDBC driver for it is called *Connector/J*. For this example we've used MySQL 4.1.7 and Connector/J 3.0.15. You can download them both from <http://www.mysql.com>.

2.1. Creating a Database and User

We'll assume that you've already installed MySQL and that you have it running and are familiar with the basics. Run the `mysql` client program from the command line so we can execute some administration commands. You should make sure that you are connected as a user with sufficient privileges (e.g. by specifying the `-u root` option to run as the MySQL root user).

First create a database called `jboss` within MySQL for use by JBoss.

```
mysql> CREATE DATABASE jboss;
Query OK, 1 row affected (0.05 sec)
```

Then check that it has been created.

```
mysql> SHOW DATABASES;
+-----+
| Database |
+-----+
| jboss    |
| mysql    |
| test     |
+-----+
3 rows in set (0.00 sec)
```

Next, create a user called `jboss` with password `password` to access the database.

```
mysql> GRANT ALL PRIVILEGES ON jboss.* TO jboss@localhost IDENTIFIED BY
'password';
Query OK, 0 rows affected (0.06 sec)
```

Again, you can check that everything has gone smoothly.

```
mysql> select User,Host,Password from mysql.User;
+-----+-----+-----+
| User  | Host      | Password          |
+-----+-----+-----+
| root  | localhost |                   |
| root  | %         |                   |
|       | localhost |                   |
|       | %         |                   |
| jboss | localhost | 5d2e19393cc5ef67 |
+-----+-----+-----+
5 rows in set (0.02 sec)
```

2.2. Installing the JDBC Driver and Deploying the DataSource

To make the JDBC driver classes available to JBoss, copy the file `mysql-connector-java-3.0.15-ga-bin.jar` from the Connector/J distribution to the `lib` directory in the `default` server configuration (assuming that is the configuration you're running, of course). Then create a file in the `deploy` directory called `mysql-ds.xml` with the following datasource configuration. The database user name and password corresponds the MySQL user we created in the previous section.

```
<datasources>
  <local-tx-datasource>
```

```

    <jndi-name>MySQLDS</jndi-name>
    <connection-url>jdbc:mysql://localhost:3306/jboss</connection-url>
    <driver-class>com.mysql.jdbc.Driver</driver-class>
    <user-name>jboss</user-name>
    <password>password</password>
  </local-tx-datasource>
</datasources>

```

Because we have added a new JAR file to the `lib` directory, you will need to make sure that the server is able to find the MySQL driver classes.

2.3. Testing the MySQL DataSource

We'll use the CMP roster application to test the new database connection. In order to use MySQL in our application, we'll need to set the `datasource` name and type-mapping in the `jbosscmp-jdbc.xml` file in the `dd/team` directory of the CMP roster application. Edit the file and add the following `datasource` and `datasource-mapping` elements to the `defaults` element.

```

<jbosscmp-jdbc>
  <defaults>
    <datasource>java:/MySQLDS</datasource>
    <datasource-mapping>mysql</datasource-mapping>
  </defaults>

  <enterprise-beans>
  ...
  </enterprise-beans>
</jbosscmp-jdbc>

```

After restarting JBoss, you should be able to deploy the application and see the tables being created. The tables should be visible from the MySQL client.

```

mysql> show tables;
+-----+
| Tables_in_jboss |
+-----+
| LeagueBean      |
| PlayerBean      |
| PlayerBean_teams_TeamBean_players |
| TeamBean        |
+-----+
4 rows in set (0.00 sec)

```

You can see the JMS persistence tables in there too since we're using MySQL as the default `datasource`.

3. Setting up an XADataSource with Oracle 9i

Oracle is one of the main players in the commercial database field and most readers will probably have come across it at some point. You can download it freely for non-commercial purposes from <http://www.oracle.com>

Installing and configuring Oracle is not for the faint of heart. It isn't really just a simple database, but it is heavy on extra features and technologies which you may not actually want (another Apache web server, multiple JDKs, Orbs etc.) but which are usually installed anyway. So we'll assume you already have an Oracle installation available. For this example, we've used Oracle 10g.

3.1. Padding Xid Values for Oracle Compatibility

If you look in the `jboss-service.xml` file in the `default/conf` directory, you'll find the following service MBean.

```
<!-- The configurable Xid factory. For use with Oracle, set pad to true -->
<mbean code="org.jboss.tm.XidFactory"
       name="jboss:service=XidFactory">
  <!--attribute name="Pad">true</attribute-->
</mbean>
```

The transaction service uses this to create XA transactions identifiers. The comment explains the situation: for use with Oracle you have to include the line which sets the attribute `Pad` to `true`. This activates padding the identifiers out to their maximum length of 64 bytes. Remember that you'll have to restart JBoss for this change to be put into effect, but wait until you've installed the JDBC driver classes which we'll talk about next.

3.2. Installing the JDBC Driver and Deploying the DataSource

The Oracle JDBC drivers can be found in the directory `$ORACLE_HOME/jdbc/lib`. Older versions, which may be more familiar to some users, had rather uninformative names like `classes12.zip` but at the time of writing the latest driver version can be found in the file `ojdbc14.jar`. There is also a debug version of the classes with `_g` appended to the name which may be useful if you run into problems. Again, you should copy one of these to the `lib` directory of the JBoss `default` configuration. The basic driver class you would use for the non-XA setup is called `oracle.jdbc.driver.OracleDriver`. The `XADataSource` class, which we'll use here, is called `oracle.jdbc.xa.client.OracleXADataSource`.

For the configuration file, make a copy of the `oracle-xa-ds.xml` example file and edit it to set the correct URL, username and password.

```
<datasources>
  <xa-datasource>
    <jndi-name>XAOracleDS</jndi-name>
    <track-connection-by-tx>true</track-connection-by-tx>
    <isSameRM-override-value>false</isSameRM-override-value>
    <xa-datasource-class>oracle.jdbc.xa.client.OracleXADataSource</xa-datasource-class>
    <xa-datasource-property name="URL">
      jdbc:oracle:thin:@monkeymachine:1521:jboss
```

```

        </xa-datasource-property>
        <xa-datasource-property name="User">jboss</xa-datasource-property>
        <xa-datasource-property
name="Password">password</xa-datasource-property>
        <exception-sorter-class-name>
            org.jboss.resource.adapter.jdbc.vendor.OracleExceptionSorter
        </exception-sorter-class-name>
        <no-tx-separate-pools/>
    </xa-datasource>

    <mbean
code="org.jboss.resource.adapter.jdbc.vender.oracle.OracleXAExceptionFormatter"
        name="jboss.jca:service=OracleXAExceptionFormatter">
        <depends optional-attribute-name="TransactionManagerService">
            jboss:service=TransactionManager
        </depends>
    </mbean>
</datasources>

```

We've used the Oracle thin (pure java) driver here and assumed the database is running on the host `monkeymachine` and that the database name (or SID in Oracle terminology) is `jboss`. We've also assumed that you've created a user `jboss` with all the sufficient privileges. You can just use `dba` privileges for this example.

```

SQL> connect / as sysdba
Connected.
SQL> create user jboss identified by password;
User created.
SQL> grant dba to jboss;
Grant succeeded.

```

Now copy the file to the `deploy` directory. You should get the following output.

```

11:33:45,174 INFO [WrapperDataSourceService] Bound connection factory for
resource adapter
for ConnectionManager 'jboss.jca:name=XAOracleDS,service=DataSourceBinding
to JNDI name
'java:XAOracleDS'

```

If you use the `JNDIView` service from the JMX console as before, you should see the name `java:/XAOracleDS` listed.

3.3. Testing the Oracle DataSource

Again we'll use the CMP example to test out the new database connection. The `jbosscmp-jdbc.xml` file should contain the following.

```

<jbosscmp-jdbc>
  <defaults>
    <datasource>java:/XAOracleDS</datasource>
  </defaults>
</jbosscmp-jdbc>

```

```
<datasource-mapping>Oracle9i</datasource-mapping>
</defaults>
</jbosscomp-jdbc>
```

There are other Oracle type-mappings available too. If you're using an earlier version, have a look in the `conf/standardjbosscomp-jdbc.xml` file to find the correct name

Deploy the application as before, check the output for errors and then check that the tables have been created using Oracle SQLPlus again from the command line.

```
SQL> select table_name from user_tables;

TABLE_NAME
-----
TEAMBEAN
LEAGUEBEAN
PLAYERBEAN
PLAYERBEAN_TEAMS_TEAM_1OFLZV8
```

Appendix A. Further Information

Sources

For a longer introduction to JBoss, see *JBoss: A Developer's Notebook*. (O'Reilly, 2005. Norman Richards, Sam Griffith).

For more comprehensive JBoss documentation covering advanced JBoss topics, refer to the manuals available online at <http://www.redhat.com/docs/manuals/jboss>.

For general EJB instruction, with thorough JBoss coverage, see *Enterprise JavaBeans, 4th Edition*. (O'Reilly, 2004. Richard Monson-Haefel, Bill Burke, Sacha Labourey)

To learn more about Hibernate, see *Java Persistence With Hibernate*. (Manning, 2007. Christian Bauer, Gavin King)

For complete coverage of the JBoss Seam framework, we recommend *JBoss Seam: Simplicity And Power Beyond Java EE*. (Prentice Hall, 2007. Michael Yuan, Thomas Heute).

