Oracle® WebLogic Server

Developing Applications with WebLogic Server

10*g* Release 3 (10.3)

July 2008

ORACLE®

Contents

Overview of WebLogic Server Application Development

Using Ant Tasks to Configure and Use a WebLogic Server Domain

Creating a Split Development Directory Environment

Building Applications in a Split Development Directory

Deploying and Packaging from a Split Development Directory

Developing Applications for Production Redeployment

Using Java EE Annotations and Dependency Injection

Understanding WebLogic Server Application Classloading

Creating Shared Java EE Libraries and Optional Packages

Programming Application Life Cycle Events

Programming Context Propagation

Programming JavaMail with WebLogic Server

Threading and Clustering Topics

Enterprise Application Deployment Descriptor Elements

wldeploy Ant Task Reference

Spring Applications Reference

Overview of WebLogic Server Application Development

The following sections provide an overview of WebLogic Server® applications and basic concepts.

Document Scope and Audience

This document is written for application developers who want to build WebLogic Server e-commerce applications using the Java Platform, Enterprise Edition 5 from Sun Microsystems. It is assumed that readers know Web technologies, object-oriented programming techniques, and the Java programming language.

WebLogic Server applications are created by Java programmers, Web designers, and application assemblers. Programmers and designers create modules that implement the business and presentation logic for the application. Application assemblers assemble the modules into applications that are ready to deploy on WebLogic Server.

WebLogic Server and the Java EE Platform

WebLogic Server implements Java Platform, Enterprise Edition (Java EE) Version 5.0 technologies. Java EE is the standard platform for developing multi-tier Enterprise applications based on the Java programming language. The technologies that make up Java EE were developed collaboratively by Sun Microsystems and other software vendors.

An important aspect of the Java EE programming model is the introduction of metadata annotations. Annotations simplify the application development process by allowing a developer to specify within the Java class itself how the application component behaves in the container, requests for dependency injection, and so on. Annotations are an alternative to deployment descriptors that were required by older versions of Enterprise applications (J2EE 1.4 and earlier).

According to Sun, "the focus in Java EE 5 is ease of development. With Java EE 5, there is less code to write – much of the boilerplate code has been removed, defaults are used whenever possible, and annotations are used extensively to reduce the need for deployment descriptors."

- EJB 3.0 makes it much easier to program an EJB, in particular by reducing the number of required programming artifacts and introducing a set of EJB-specific metadata annotations that make programming the bean file easier and more intuitive. Another goal of EJB 3.0 is to standardize the persistence framework and reduce the complexity of the entity bean programming model and object-relational (O/R) mapping model. WebLogic Server continues to support Version 2.1 of the EJB specification.

- Java EE 5 includes simplified Web Services support and the latest web services APIs, making it an ideal implementation platform for Service-Oriented Architectures (SOA).

- Constructing web applications is made easier with JavaServer Faces (JSF) technology and the JSP Standard Tag Library (JSTL). Java EE 5 supports rich thin-client technologies such as AJAX, for building applications for Web 2.0.

WebLogic Server Java EE applications are based on standardized, modular components. WebLogic Server provides a complete set of services for those modules and handles many details of application behavior automatically, without requiring programming. Java EE defines module behaviors and packaging in a generic, portable way, postponing run-time configuration until the module is actually deployed on an application server.

Java EE includes deployment specifications for Web applications, EJB modules, Web Services, Enterprise applications, client applications, and connectors. Java EE does not specify *how* an application is deployed on the target server—only how a standard module or application is packaged. For each module type, the specifications define the files required and their location in the directory structure.

Note: Because Java EE is backward compatible, you can still run Java EE applications on WebLogic Server versions 8.1 and higher.

Java is platform independent, so you can edit and compile code on any platform, and test your applications on development WebLogic Servers running on other platforms. For example, it is common to develop WebLogic Server applications on a PC running Windows or Linux, regardless of the platform where the application is ultimately deployed.

For more information, refer to the Java EE specification at:
`http://java.sun.com/javaee/5/docs/api/`

Overview of Java EE Applications and Modules

A WebLogic Server Java EE application consists of one of the following modules or applications running on WebLogic Server:

- Web application modules—HTML pages, servlets, JavaServer Pages, and related files. See "Web Application Modules" on page 1-4.

- Enterprise Java Beans (EJB) modules—entity beans, session beans, and message-driven beans. See "Enterprise JavaBean Modules" on page 1-5.

- Connector modules—resource adapters. See "Connector Modules" on page 1-6.

- Enterprise applications—Web application modules, EJB modules, resource adapters and Web Services packaged into an application. See "Enterprise Applications" on page 1-7.

- Web Services—See "WebLogic Web Services" on page 1-8.

A WebLogic application can also include the following WebLogic-specific modules:

- JDBC and JMS modules—See "JMS and JDBC Modules" on page 1-9.

- WebLogic Diagnostic FrameWork (WLDF) modules—See "WebLogic Diagnostic Framework Modules" on page 1-10.

Web Application Modules

A Web application on WebLogic Server includes the following files:

- At least one servlet or JSP, along with any helper classes.

- Optionally, a web.xml deployment descriptor, a Java EE standard XML document that describes the contents of a WAR file.

- Optionally, a weblogic.xml deployment descriptor, an XML document containing WebLogic Server-specific elements for Web applications.

- A Web application can also include HTML and XML pages with supporting files such as images and multimedia files.

Servlets

Servlets are Java classes that execute in WebLogic Server, accept a request from a client, process it, and optionally return a response to the client. An HttpServlet is most often used to generate dynamic Web pages in response to Web browser requests.

JavaServer Pages

JavaServer Pages (JSPs) are Web pages coded with an extended HTML that makes it possible to embed Java code in a Web page. JSPs can call custom Java classes, known as tag libraries, using HTML-like tags. The appc compiler compiles JSPs and translates them into servlets. WebLogic Server automatically compiles JSPs if the servlet class file is not present or is older than the JSP source file. See "Building Modules and Applications Using wlappc" on page 4-4.

You can also precompile JSPs and package the servlet class in a Web Application (WAR) file to avoid compiling in the server. Servlets and JSPs may require additional helper classes that must also be deployed with the Web application.

More Information on Web Application Modules

See the following documentation:

- "Organizing Java EE Components in a Split Development Directory" on page 3-6.

- Developing Web Applications, Servlets, and JSPs for WebLogic Server

- Programming JSP Tag Extensions

Enterprise JavaBean Modules

Enterprise JavaBeans (EJBs) beans are server-side Java modules that implement a business task or entity and are written according to the EJB specification. There are three types of EJBs: session beans, entity beans, and message-driven beans.

Enterprise JavaBeans (EJB) 3.0 is a Java EE 5 technology for the development and deployment of component-based business applications. Although EJB 2.X is a powerful and useful technology, the programming model was complex and confusing, requiring the creation of multiple Java files and deployment descriptors for even the simplest EJB. This complexity hindered the wide adoption of EJBs.

Therefore, one of the central goals of version 3.0 of the EJB specification is to make it easier to program an EJB, in particular by reducing the number of required programming artifacts and introducing a set of EJB-specific metadata annotations that make programming the bean file easier and more intuitive. Another goal of the EJB 3.0 specification was to standardize the persistence framework and reduce the complexity of the entity bean programming model and object-relational (O/R) mapping model.

For more information on Enterprise JavaBeans 3.0, see Understanding WebLogic Enterprise JavaBeans 3.0.

For more information on Enterprise JavaBeans 2.X, see Understanding WebLogic Enterprise JavaBeans.

EJB Overview

Session beans execute a particular business task on behalf of a single client during a single session. Session beans can be stateful or stateless, but are not persistent; when a client finishes with a session bean, the bean goes away.

Entity beans represent business objects in a data store, usually a relational database system. Persistence—loading and saving data—can be bean-managed or container-managed. More than just an in-memory representation of a data object, entity beans have methods that model the behaviors of the business objects they represent. Entity beans can be accessed concurrently by multiple clients and they are persistent by definition.

The container creates an instance of the message-driven bean or it assigns one from a pool to process the message. When the message is received in the JMS destination, the message-driven bean assigns an instance of itself from a pool to process the message. Message-driven beans are not associated with any client. They simply handle messages as they arrive.

EJBs and WebLogic Server

Java EE cleanly separates the development and deployment roles to ensure that modules are portable between EJB servers that support the EJB specification. Deploying an EJB in WebLogic Server requires running the WebLogic Server `appc` compiler to generate classes that enforce the EJB security, transaction, and life cycle policies. See "Building Modules and Applications Using wlappc" on page 4-4.

The Java EE-specified deployment descriptor, `ejb-jar.xml`, describes the enterprise beans packaged in an EJB application. It defines the beans' types, names, and the names of their home and remote interfaces and implementation classes. The `ejb-jar.xml` deployment descriptor defines security roles for the beans, and transactional behaviors for the beans' methods.

Additional deployment descriptors provide WebLogic-specific deployment information. A `weblogic-cmp-rdbms-jar.xml` deployment descriptor unique to container-managed entity beans maps a bean to tables in a database. The `weblogic-ejb-jar.xml` deployment descriptor supplies additional information specific to the WebLogic Server environment, such as JNDI bind names, clustering, and cache configuration.

Connector Modules

Connectors (also known as resource adapters) contain the Java, and if necessary, the native modules required to interact with an Enterprise Information System (EIS). A resource adapter deployed to the WebLogic Server environment enables Java EE applications to access a remote EIS. WebLogic Server application developers can use HTTP servlets, JavaServer Pages (JSPs), Enterprise Java Beans (EJBs), and other APIs to develop integrated applications that use the EIS data and business logic.

To deploy a resource adapter to WebLogic Server, you must first create and configure WebLogic Server-specific deployment descriptor, `weblogic-ra.xml` file, and add this to the deployment directory. Resource adapters can be deployed to WebLogic Server as stand-alone modules or as part of an Enterprise application. See "Enterprise Applications" on page 1-7.

For more information on connectors, see *Programming WebLogic Resource Adapters.*

Enterprise Applications

An Enterprise application consists of one or more Web application modules, EJB modules, and resource adapters. It might also include a client application. An Enterprise application can be optionally defined by an `application.xml` file, which was the standard J2EE deployment descriptor for Enterprise applications.

Java EE Programming Model

An important aspect of the Java EE programming model is the introduction of metadata annotations. Annotations simplify the application development process by allowing a developer to specify within the Java class itself how the application behaves in the container, requests for dependency injection, and so on. Annotations are an alternative to deployment descriptors that were required by older versions of Enterprise applications (1.4 and earlier).

With Java EE annotations, the standard `application.xml` and `web.xml` deployment descriptors are optional. The Java EE programming model uses the JDK 5.0 annotations feature for Web containers, such as EJBs, servlets, Web applications, and JSPs. See Chapter 7, "Using Java EE Annotations and Dependency Injection."

If the application includes WebLogic Server-specific extensions, the application is further defined by a `weblogic-application.xml` file. Enterprise Applications that include a client module will also have a `client-application.xml` deployment descriptor and a WebLogic run-time client application deployment descriptor. See Appendix A, "Enterprise Application Deployment Descriptor Elements."

Packaging and Deployment Overview

For both production and development purposes, Oracle recommends that you package and deploy even stand-alone Web applications, EJBs, and resource adapters as part of an Enterprise application. Doing so allows you to take advantage of Oracle's split development directory structure, which greatly facilities application development. See Chapter 3, "Creating a Split Development Directory Environment."

An Enterprise application consists of Web application modules, EJB modules, and resource adapters. It can be packaged as follows:

- For development purposes, Oracle recommends the WebLogic split development directory structure. Rather than having a single archived EAR file or an exploded EAR directory structure, the split development directory has two parallel directories that separate source files and output files. This directory structure is optimized for development on a single

WebLogic Server instance. See Chapter 3, "Creating a Split Development Directory Environment." Oracle provides the `wlpackage` Ant task, which allows you to create an EAR without having to use the JAR utility; this is exclusively for the split development directory structure. See "Packaging Applications Using wlpackage" on page 5-1.

- For development purposes, Oracle further recommends that you package stand-alone Web applications and Enterprise JavaBeans (EJBs) as part of an Enterprise application, so that you can take advantage of the split development directory structure. See "Organizing Java EE Components in a Split Development Directory" on page 3-6.

- For production purposes, Oracle recommends the exploded (unarchived) directory format. This format enables you to update files without having to redeploy the application. To update an archived file, you must unarchive the file, update it, then rearchive and redeploy it.

- You can choose to package your application as a JAR archived file using the `jar` utility with an `.ear` extension. Archived files are easier to distribute and take up less space. An EAR file contains all of the JAR, WAR, and RAR module archive files for an application and an XML descriptor that describes the bundled modules. See "Packaging Applications Using wlpackage" on page 5-1.

The optional `META-INF/application.xml` deployment descriptor contains an element for each Web application, EJB, and connector module, as well as additional elements to describe security roles and application resources such as databases. If this descriptor is present the WebLogic deployer picks the list of modules from this descriptor. However if this descriptor is not present, the container guesses the modules from the annotations defined on the POJO (plain-old-Java-object) classes. See Appendix A, "Enterprise Application Deployment Descriptor Elements."

WebLogic Web Services

Web Services can be shared by and used as modules of distributed Web-based applications. They commonly interface with existing back-end applications, such as customer relationship management systems, order-processing systems, and so on. Web Services can reside on different computers and can be implemented by vastly different technologies, but they are packaged and transported using standard Web protocols, such as HTTP, thus making them easily accessible by any user on the Web.

A Web Service consists of the following modules, at a minimum:

- A Web Service implementation hosted by a server on the Web. WebLogic Web Services are hosted by WebLogic Server. A Web Service module may include either Java classes or

EJBs that implement the Web Service. Web Services are packaged either as Web Application archives (WARs) or EJB modules (JARs), depending on the implementation.

- A standard for transmitting data and Web Service invocation calls between the Web Service and the user of the Web Service. WebLogic Web Services use Simple Object Access Protocol (SOAP) 1.1 as the message format and HTTP as the connection protocol.

- A standard for describing the Web Service to clients so they can invoke it. WebLogic Web Services use Web Services Description Language (WSDL) 1.1, an XML-based specification, to describe themselves.

- A standard for clients to invoke Web services—JAX-WS or JAX-RPC. See Getting Started With WebLogic Web Services Using JAX-WS or Getting Started With WebLogic Web Services Using JAX-RPC, respectively.

- A standard for finding and registering the Web Service (UDDI).

For more information about WebLogic Web Services and the standards that are supported, see Web Services for WebLogic Server.

JMS and JDBC Modules

JMS and JDBC configurations are stored as modules, defined by an XML file that conforms to the `weblogic-jms.xsd` and `jdbc-data-source.xsd` schema, respectively. These modules are similar to standard Java EE modules. An administrator can create and manage JMS and JDBC modules as global system resources, as modules packaged with a Java EE application (as a packaged resource), or as standalone modules that can be made globally available.

With modular deployment of JMS and JDBC resources, you can migrate your application and the required JMS or JDBC configuration from environment to environment, such as from a testing environment to a production environment, without opening an enterprise application file (such as an EAR file) or a JMS or JDBC standalone module, and without extensive manual JMS or JDBC reconfiguration.

Application developers create application modules in an enterprise-level IDE or another development tool that supports editing of XML files, then package the JMS or JDBC modules with an application and pass the application to a WebLogic Administrator to deploy.

For more information, see:

- Configuring JMS Application Modules for Deployment

- Configuring JDBC Application Modules for Deployment

WebLogic Diagnostic Framework Modules

The WebLogic Diagnostic Framework (WLDF) provides features for generating, gathering, analyzing, and persisting diagnostic data from WebLogic Server instances and from applications deployed to server instances. For server-scoped diagnostics, some WLDF features are configured as part of the configuration for the domain. Other features are configured as system resource descriptors that can be targeted to servers (or clusters). For application-scoped diagnostics, diagnostic features are configured as resource descriptors for the application.

Application-scoped instrumentation is configured and deployed as a diagnostic module, which is similar to a diagnostic system module. However, an application module is configured in an XML configuration file named `weblogic-diagnostics.xml` which is packaged with the application archive.

For detailed instructions for configuring instrumentation for applications, see Configuring Application-Scoped Instrumentation.

Using an External Diagnostics Descriptor

WLS also supports the use of an external diagnostics descriptor so you can integrate diagnostic functionality into an application that has not imported diagnostic descriptors. This feature supports the deployment view and deployment of an application or a module, detecting the presence of an external diagnostics descriptor if the descriptor is defined in your deployment plan (`plan.xml`).

Defining an External Diagnostics Descriptor

First, define the diagnostic descriptor as external and configure its URI in the `plan.xml` file. For example:

```
<module-override>
  <module-name>reviewService.ear</module-name>
  <module-type>ear</module-type>
  </module-descriptor>
  <module-descriptor external="true">
  <root-element>wldf-resource</root-element>
  <uri>META-INF/weblogic-diagnostics.xml</uri>
  ...
  ...
</module-override>
<config-root>D:\plan</config-root>
```

Then place the external diagnostic descriptor file under the URI. Using the example above, you would place the descriptor file under `d:\plan\ META-INF`.

XML Deployment Descriptors

A *deployment configuration* refers to the process of defining the deployment descriptor values required to deploy an Enterprise application to a particular WebLogic Server domain. The deployment configuration for an application or module is stored in three types of XML document: Java EE deployment descriptors, WebLogic Server descriptors, and WebLogic Server deployment plans. This section describes the Java EE and WebLogic-specific deployment descriptors. See "Deployment Plans" on page 1-17 for information on deployment plans.

The Java EE programming model uses the JDK 5.0 annotations feature for Web containers, such as EJBs, servlets, Web applications, and JSPs. Annotations simplify the application development process by allowing a developer to specify within the Java class itself how the component behaves in the container, requests for dependency injection, and so on. Annotations are an alternative to deployment descriptors that were required by older versions of Web Applications (2.4 and earlier), Enterprise applications, and (1.4 and earlier), and Enterprise JavaBeans (2.*x* and earlier). See Chapter 7, "Using Java EE Annotations and Dependency Injection."

However, Enterprise applications fully support the use of deployment descriptors, even though the standard J2EE ones are not required. For example, you may prefer to use the old EJB 2.*x* programming model, or might want to allow further customizing of the EJB at a later development or deployment stage; in these cases you can create the standard deployment descriptors in addition to, or instead of, the metadata annotations.

Modules and applications have deployment descriptors—XML documents—that describe the contents of the directory or JAR file. Deployment descriptors are text documents formatted with XML tags. The Java EE specifications define standard, portable deployment descriptors for Java EE modules and applications. Oracle defines additional WebLogic-specific deployment descriptors for deploying a module or application in the WebLogic Server environment.

Table 1-1 lists the types of modules and applications and their Java EE-standard and WebLogic-specific deployment descriptors.

Note: The XML schemas for the WebLogic deployment descriptors listed in the following table include elements from the http://www.bea.com/ns/weblogic/weblogic-javaee/1.0/weblogic-javaee.xsd schema,

which describes common elements shared among all WebLogic-specific deployment descriptors.

Table 1-1 Java EE and WebLogic Deployment Descriptors

Module or Application	Scope	Deployment Descriptors
Web Application	Java EE	`web.xml` See the Sun Microsystems Servlet 2.5 Schema.
	WebLogic	`weblogic.xml` Schema: See weblogic.xml Deployment Descriptor Elements in *Developing Web Applications for WebLogic Server.*
Enterprise Bean 3.0	Java EE	`ejb-jar.xml` See the Sun Microsystems EJB 3.0 Schema.

	WebLogic	`weblogic-ejb-jar.xml`
		Schema: http://www.bea.com/ns/weblogic/weblogic-ejb-jar/1.0/weblogic-ejb-jar.xsd
		`weblogic-rdbms-jar.xml`
		Schema: http://www.bea.com/ns/weblogic/weblogic-rdbms-jar/1.0/weblogic-rdbms-jar.xsd
		`persistence-configuration.xml`
		Schema: http://www.bea.com/ns/weblogic/persistence-configuration/1.0/persistence-configuration.xsd
		See *Programming WebLogic Enterprise JavaBeans, Version 3.0.*

Table 1-1 Java EE and WebLogic Deployment Descriptors

Module or Application	Scope	Deployment Descriptors
Enterprise Bean 2.1	J2EE	`ejb-jar.xml`
		See the Sun Microsystems EJB 2.1 Schema.
	WebLogic	`weblogic-ejb-jar.xml`
		Schema: http://www.bea.com/ns/weblogic/10.0/weblogic-ejb-jar.xsd
		See "The weblogic-ejb-jar.xml Deployment Descriptor" in *Programming WebLogic Enterprise JavaBeans*.
		`weblogic-cmp-rdbms-jar.xml`
		Schema: http://www.bea.com/ns/weblogic/920/weblogic-rdbms20-persistence.xsd
		See "The weblogic-cmp-rdbms-jar.xml Deployment Descriptor" in *Programming WebLogic Enterprise JavaBeans*.
Web Services	Java EE	`webservices.xml`
		See the Sun Microsystems Web Services 1.2 Schema.
	WebLogic	`weblogic-webservices.xml`
		Schema: http://www.bea.com/ns/weblogic/weblogic-webservices/1.0/weblogic-webservices.xsd
		`weblogic-wsee-clientHandlerChain.xml`
		Schema: http://www.bea.com/ns/weblogic/weblogic-wsee-clientHandlerChain1.0/weblogic-wsee-clientHandlerChain.xsd
		`webservice-policy-ref.xml`
		Schema: http://www.bea.com/ns/weblogic/webservice-policy-ref/1.0/webservice-policy-ref.xsd
		`weblogic-wsee-standaloneclient.xml`
		Schema: http://www.bea.com/ns/weblogic/weblogic-wsee-standaloneclient/1.0/weblogic-wsee-standaloneclient.xsd
		See the WebLogic Web Service Deployment Descriptor Element Reference in *WebLogic Web Services Reference*.

Table 1-1 Java EE and WebLogic Deployment Descriptors

Module or Application	Scope	Deployment Descriptors
Resource Adapter	Java EE	`ra.xml` See the Sun Microsystems Connector 1.5 Schema.
	WebLogic	`weblogic-ra.xml` Schema: http://www.bea.com/ns/weblogic/weblogic-connector/1.0/weblogic-connector.xsd See weblogic-ra.xml Schema in *Programming WebLogic Resource Adapters*.
Enterprise Application	Java EE	`application.xml` See the Sun Microsystems Application 5 Schema.
	WebLogic	`weblogic-application.xml` Schema: http://www.bea.com/ns/weblogic/weblogic-application/1.0/weblogic-application.xsd See "weblogic-application.xml Deployment Descriptor Elements" on page A-1.
Client Application	Java EE	`application-client.xml` See the Sun Microsystems Application Client 5 Schema.
	WebLogic	`application-client.xml` Schema: http://www.bea.com/ns/weblogic/weblogic-application-client/1.0/weblogic-application-client.xsd See Developing a J2EE Application Client (Thin Client) in *Programming Stand-alone Clients*.
JMS Module	WebLogic	*FileName*-`jms.xml`, where *FileName* can be anything you want. Schema: http://www.bea.com/ns/weblogic/weblogic-jms/1.0/weblogic-jms.xsd See Configuring JMS Application Modules for Deployment in *Configuring and Managing WebLogic JMS*.

Table 1-1 Java EE and WebLogic Deployment Descriptors

Module or Application	Scope	Deployment Descriptors
JDBC Module	WebLogic	`FileName-jdbc.xml`, where `FileName` can be anything you want.
		Schema: http://www.bea.com/ns/weblogic/jdbc-data-source/1.0/jdbc-data-source.xsd
		See Configuring JDBC Application Modules for Deployment in *Configuring and Managing WebLogic JDBC*.
Deployment Plan	WebLogic	`plan.xml`
		Schema: http://www.bea.com/ns/weblogic/deployment-plan/1.0/deployment-plan.xsd
		See Understanding WebLogic Server Deployment in *Deploying Applications to WebLogic Server*.
WLDF Module	WebLogic	`weblogic-diagnostics.xml`
		Schema: http://www.bea.com/ns/weblogic/weblogic-diagnostics/1.1/weblogic-diagnostics.xsd
		See Deploying WLDF Application Modules in *Configuring and Using the WebLogic Diagnostics Framework*.

When you package a module or application, you create a directory to hold the deployment descriptors—`WEB-INF` or `META-INF`—and then create the XML deployment descriptors in that directory.

Automatically Generating Deployment Descriptors

WebLogic Server provides a variety of tools for automatically generating deployment descriptors. These are discussed in the sections that follow.

EJBGen

EJBGen is an Enterprise JavaBeans 2.*x* code generator or command-line tool that uses Javadoc markup to generate EJB deployment descriptor files. You annotate your Bean class file with Javadoc tags and then use EJBGen to generate the Remote and Home classes and the deployment

descriptor files for an EJB application, reducing to a single file you need to edit and maintain your EJB .java and descriptor files. See "EJBGen Reference" in *Programming WebLogic Enterprise JavaBeans*.

Java-based Command-line Utilities

WebLogic Server includes a set of Java-based command-line utilities that automatically generate both standard Java EE and WebLogic-specific deployment descriptors for Web applications and Enterprise Applications.

These command-line utilities examine the classes you have assembled in a staging directory and build the appropriate deployment descriptors based on the servlet classes, and so on. These utilities include:

- `java weblogic.marathon.ddinit.EarInit` — automatically generates the deployment descriptors for Enterprise applications.

- `java weblogic.marathon.ddinit.WebInit` — automatically generates the deployment descriptors for Web applications.

For an example of DDInit, assume that you have created a directory called `c:\stage` that contains the JSP files and other objects that make up a Web application but you have not yet created the `web.xml` and `weblogic.xml` deployment descriptors. To automatically generate them, execute the following command:

```
prompt> java weblogic.marathon.ddinit.WebInit c:\stage
```

The utility generates the `web.xml` and `weblogic.xml` deployment descriptors and places them in the `WEB-INF` directory, which DDInit will create if it does not already exist.

Upgrading Deployment Descriptors From Previous Releases of J2EE and WebLogic Server

So that your applications can take advantage of the features in the current Java EE specification and release of WebLogic Server, Oracle recommends that you always upgrade deployment descriptors when you migrate applications to a new release of WebLogic Server.

To upgrade the deployment descriptors in your J2EE applications and modules, first use the `weblogic.DDConverter` tool to generate the upgraded descriptors into a temporary directory. Once you have inspected the upgraded deployment descriptors to ensure that they are correct, repackage your J2EE module archive or exploded directory with the new deployment descriptor files.

Invoke `weblogic.DDConverter` with the following command:

```
prompt> java weblogic.DDConverter [options] archive_file_or_directory
```

where *archive_file_or_directory* refers to the archive file (EAR, WAR, JAR, or RAR) or exploded directory of your Enterprise application, Web application, EJB, or resource adapter.

The following table describes the `weblogic.DDConverter` command options.

Table 1-2 `weblogic.DDConverter` **Command Options**

Option	Description
-d <dir>	Specifies the directory to which descriptors are written.
-help	Prints the standard usage message.
-quiet	Turns off output messages except error messages.
-verbose	Turns on additional output used for debugging.

The following example shows how to use the `weblogic.DDConverter` command to generate upgraded deployment descriptors for the `my.ear` Enterprise application into the subdirectory `tempdir` in the current directory:

```
prompt> java weblogic.DDConverter -d tempdir my.ear
```

Deployment Plans

A *deployment plan* is an XML document that defines an application's WebLogic Server deployment configuration for a specific WebLogic Server environment. A deployment plan resides outside of an application's archive file, and can apply changes to deployment properties stored in the application's existing WebLogic Server deployment descriptors. Administrators use deployment plans to easily change an application's WebLogic Server configuration for a specific environment *without* modifying existing Java EE or WebLogic-specific deployment descriptors. Multiple deployment plans can be used to reconfigure a single application for deployment to multiple, differing WebLogic Server environments.

After programmers have finished programming an application, they export its deployment configuration to create a custom deployment plan that administrators later use for deploying the

application into new WebLogic Server environments. Programmers distribute both the application deployment files and the custom deployment plan to deployers (for example, testing, staging, or production administrators) who use the deployment plan as a blueprint for configuring the application for their environment.

WebLogic Server provides the following tools to help programmers export an application's deployment configuration:

- `weblogic.PlanGenerator` creates a template deployment plan with null variables for selected categories of WebLogic Server deployment descriptors. This tool is recommended if you are beginning the export process and you want to create a template deployment plan with null variables for an entire class of deployment descriptors.

- The Administration Console updates or creates new deployment plans as necessary when you change configuration properties for an installed application. You can use the Administration Console to generate a new deployment plan or to add or override variables in an existing plan. The Administration Console provides greater flexibility than `weblogic.PlanGenerator`, because it allows you to interactively add or edit individual deployment descriptor properties in the plan, rather than export entire categories of descriptor properties.

For complete and detailed information about creating and using deployment plans, see:

- Understanding WebLogic Server Deployment

- Exporting an Application for Deployment to New Environments

- Understanding WebLogic Server Deployment Plans

Development Software

This section reviews required and optional tools for developing WebLogic Server applications.

Apache Ant

The preferred Oracle method for building applications with WebLogic Server is Apache Ant. Ant is a Java-based build tool. One of the benefits of Ant is that is it is extended with Java classes, rather than shell-based commands. Oracle provides numerous Ant extension classes to help you compile, build, deploy, and package applications using the WebLogic Server split development directory environment.

Another benefit is that Ant is a cross-platform tool. Developers write Ant build scripts in eXtensible Markup Language (XML). XML tags define the targets to build, dependencies among

targets, and tasks to execute in order to build the targets. Ant libraries are bundled with WebLogic Server to make it easier for our customers to build Java applications out of the box.

To use Ant, you must first set your environment by executing either the `setExamplesEnv.cmd` (Windows) or `setExamplesEnv.sh` (UNIX) commands located in the `WL_SERVER\samples\domains\wl_server` directory, where `WL_SERVER` is your WebLogic Server installation directory.

For a complete explanation of ant capabilities, see:
http://jakarta.apache.org/ant/manual/index.html

Note: The Apache Jakarta Web site publishes online documentation for only the most current version of Ant, which might be different from the version of Ant that is bundled with WebLogic Server. Use the following command, after setting your WebLogic environment, to determine the version of Ant bundled with WebLogic Server:

```
prompt> ant -version
```

To view the documentation for a specific version of Ant, such as the version included with WebLogic Server, download the Ant zip file from http://archive.apache.org/dist/ant/binaries/ and extract the documentation.

For more information on using Ant to compile your cross-platform scripts or using cross-platform scripts to create XML scripts that can be processed by Ant, refer to any of the WebLogic Server examples, such as
`WL_HOME/samples/server/examples/src/examples/ejb20/basic/beanManaged/build.xml`.

Also refer to the following WebLogic Server documentation on building examples using Ant:
`WL_HOME/samples/server/examples/src/examples/examples.html`.

Using A Third-Party Version of Ant

You can use your own version of Ant if the one bundled with WebLogic Server is not adequate for your purposes. To determine the version of Ant that is bundled with WebLogic Server, run the following command after setting your WebLogic environment:

```
prompt> ant -version
```

If you plan to use a different version of Ant, you can replace the appropriate JAR file in the `WL_HOME\server\lib\ant` directory with an updated version of the file (where `WL_HOME` refers to the main WebLogic installation directory, such as `c:\bea\wlserver_10.x`) or add the new file to the front of your CLASSPATH.

Changing the Ant Heap Size

By default the environment script allocates a heap size of 128 megabytes to Ant. You can increase or decrease this value for your own projects by setting the `-X` option in your local `ANT_OPTS` environment variable. For example:

```
prompt> setenv ANT_OPTS=-Xmx128m
```

If you want to set the heap size permanently, add or update the `MEM_ARGS` variable in the scripts that set your environment, start WebLogic Server, and so on, as shown in the following snippet from a Windows command script that starts a WebLogic Server instance:

```
set MEM_ARGS=-Xms32m -Xmx200m
```

See the scripts and commands in `WL_HOME`/`server`/`bin` for examples of using the `MEM_ARGS` variable.

Source Code Editor or IDE

You need a text editor to edit Java source files, configuration files, HTML or XML pages, and JavaServer Pages. An editor that gracefully handles Windows and UNIX line-ending differences is preferred, but there are no other special requirements for your editor. You can edit HTML or XML pages and JavaServer Pages with a plain text editor, or use a Web page editor such as DreamWeaver. For XML pages, you can also use an enterprise-level IDE with DTD validation or another development tool that supports editing of XML files.

Database System and JDBC Driver

Nearly all WebLogic Server applications require a database system. You can use any DBMS that you can access with a standard JDBC driver, but services such as WebLogic Java Message Service (JMS) require a supported JDBC driver for Oracle, Sybase, Informix, Microsoft SQL Server, IBM DB2, or PointBase. Refer to *Supported Configurations* to find out about supported database systems and JDBC drivers.

Web Browser

Most Java EE applications are designed to be executed by Web browser clients. WebLogic Server supports the HTTP 1.1 specification and is tested with current versions of the Firefox and Microsoft Internet Explorer browsers.

When you write requirements for your application, note which Web browser versions you will support. In your test plans, include testing plans for each supported version. Be explicit about version numbers and browser configurations. Will your application support Secure Socket Layers

(SSL) protocol? Test alternative security settings in the browser so that you can tell your users what choices you support.

If your application uses applets, it is especially important to test browser configurations you want to support because of differences in the JVMs embedded in various browsers. One solution is to require users to install the Java plug-in from Sun so that everyone has the same Java run-time version.

Third-Party Software

You can use third-party software products to enhance your WebLogic Server development environment. See *WebLogic Developer Tools Resources*, which provides developer tools information for products that support the application servers.

Note: Check with the software vendor to verify software compatibility with your platform and WebLogic Server version.

Using Ant Tasks to Configure and Use a WebLogic Server Domain

The following sections describe how to start and stop WebLogic Server instances and configure WebLogic Server domains using WebLogic Ant tasks that you can include in your development build scripts:

Overview of Configuring and Starting Domains Using Ant Tasks

WebLogic Server provides a pair of Ant tasks to help you perform common configuration tasks in a development environment. The configuration tasks enable you to start and stop WebLogic Server instances as well as create and configure WebLogic Server domains.

When combined with other WebLogic Ant tasks, you can create powerful build scripts for demonstrating or testing your application with custom domains. For example, a single Ant build script can:

- Compile your application using the `wlcompile`, `wlappc`, and Web Services Ant tasks.

- Create a new single-server domain and start the Administration Server using the `wlserver` Ant task.

- Configure the new domain with required application resources using the `wlconfig` Ant task.

- Deploy the application using the `wldeploy` Ant task.

- Automatically start a compiled client application to demonstrate or test product features.

The sections that follow describe how to use the configuration Ant tasks, `wlserver` and `wlconfig`.

Starting Servers and Creating Domains Using the wlserver Ant Task

The `wlserver` Ant task enables you to start, reboot, shutdown, or connect to a WebLogic Server instance. The server instance may already exist in a configured WebLogic Server domain, or you can create a new single-server domain for development by using the `generateconfig=true` attribute.

When you use the `wlserver` task in an Ant script, the task does not return control until the specified server is available and listening for connections. If you start up a server instance using wlserver, the server process automatically terminates after the Ant VM terminates. If you only connect to a currently-running server using the `wlserver` task, the server process keeps running after Ant completes.

The `wlserver` WebLogic Server Ant task extends the standard `java` Ant task (`org.apache.tools.ant.taskdefs.Java`). This means that all the attributes of the `java` Ant task also apply to the `wlserver` Ant task. For example, you can use the `output` and `error` attributes to specify the name of the files to which output and standard errors of the `wlserver` Ant task is written, respectively. For full documentation about the attributes of the standard java Ant task, see Java on the Apache Ant site.

Basic Steps for Using wlserver

To use the wlserver Ant task:

1. Set your environment.

   On Windows NT, execute the `setWLSEnv.cmd` command, located in the directory `WL_HOME\server\bin`, where `WL_HOME` is the top-level directory of your WebLogic Server installation.

On UNIX, execute the `setWLSEnv.sh` command, located in the directory `WL_HOME/server/bin`, where `WL_HOME` is the top-level directory of your WebLogic Server installation.

Note: The `wlserver` task is predefined in the version of Ant shipped with WebLogic Server. If you want to use the task with your own Ant installation, add the following task definition in your build file:

```
<taskdef name="wlserver"
classname="weblogic.ant.taskdefs.management.WLServer"/>
```

2. Add a call to the `wlserver` task in the build script to start, shutdown, restart, or connect to a server. See "wlserver Ant Task Reference" on page 2-4 for information about `wlserver` attributes and default behavior.

3. Execute the Ant task or tasks specified in the `build.xml` file by typing `ant` in the staging directory, optionally passing the command a target argument:

```
prompt> ant
```

Use `ant -verbose` to obtain more detailed messages from the `wlserver` task.

Sample build.xml Files for wlserver

The following shows a minimal `wlserver` target that starts a server in the current directory using all default values:

```
<target name="wlserver-default">
  <wlserver/>
</target>
```

This target connects to an existing, running server using the indicated connection parameters and username/password combination:

```
<target name="connect-server">
   <wlserver host="127.0.0.1" port="7001" username="weblogic"
password="weblogic" action="connect"/>
  </target>
```

This target starts a WebLogic Server instance configured in the `config` subdirectory:

```
<target name="start-server">
  <wlserver dir="./config" host="127.0.0.1" port="7001" action="start"/>
</target>
```

This target creates a new single-server domain in an empty directory, and starts the domain's server instance:

```
<target name="new-server">
  <delete dir="./tmp"/>
  <mkdir dir="./tmp"/>
  <wlserver dir="./tmp" host="127.0.0.1" port="7001"
  generateConfig="true" username="weblogic" password="weblogic"
action="start"/>
  </target>
```

wlserver Ant Task Reference

The following table describes the attributes of the wlserver Ant task.

Table 2-1 Attributes of the wlserver Ant Task

Attribute	Description	Data Type	Required?
policy	The path to the security policy file for the WebLogic Server domain. This attribute is used only for starting server instances.	File	No
dir	The path that holds the domain configuration (for example, c:\bea\user_projects\mydomain). By default, wlserver uses the current directory.	File	No
beahome	The path to the home directory (for example, c:\bea).	File	No
weblogichome	The path to the WebLogic Server installation directory (for example, c:\bea\wlserver_10.3).	File	No
servername	The name of the server to start, shutdown, reboot, or connect to. A WebLogic Server instance is uniquely identified by its protocol, host, and port values, so if you use this set of attributes to specify the server you want to start, shutdown or reboot, you do not need to specify its actual name using the servername attribute. The only exception is when you want to shutdown the Administration server; in this case you *must* specify this attribute. The default value for this attribute is myserver.	String	Required only when shutting down the Administration server.

Table 2-1 Attributes of the wlserver Ant Task (Continued)

Attribute	Description	Data Type	Required?
domainname	The name of the WebLogic Server domain in which the server is configured.	String	No
adminserverurl	The URL to access the Administration Server in the domain. This attribute is required if you are starting up a Managed Server in the domain.	String	Required for starting Managed Servers.
username	The username of an administrator account. If you omit both the username and password attributes, wlserver attempts to obtain the encrypted username and password values from the boot.properties file. See Boot Identity Files in the *Managing Server Startup and Shutdown* for more information on boot.properties.	String	No
password	The password of an administrator account. If you omit both the username and password attributes, wlserver attempts to obtain the encrypted username and password values from the boot.properties file. See Boot Identity Files in the *Managing Server Startup and Shutdown* for more information on boot.properties.	String	No
pkpassword	The private key password for decrypting the SSL private key file.	String	No
timeout	The maximum time, in milliseconds, that wlserver waits for a server to boot. This also specifies the maximum amount of time to wait when connecting to a running server. The default value for this attribute is 0, which means the Ant task never times out.	long	No
timeoutSeconds	The maximum time, in seconds, that wlserver waits for a server to boot. This also specifies the maximum amount of time to wait when connecting to a running server. The default value for this attribute is 0, which means the Ant task never times out.	long	No

Table 2-1 Attributes of the wlserver Ant Task (Continued)

Attribute	Description	Data Type	Required?
productionmodeenabled	Specifies whether a server instance boots in development mode or in production mode.	boolean	No
	Development mode enables a WebLogic Server instance to automatically deploy and update applications that are in the *domain_name*/autodeploy directory (where *domain_name* is the name of a WebLogic Server domain). In other words, development mode lets you use auto-deploy. Production mode disables the auto-deployment feature. See *Deploying Applications and Modules* for more information.		
	Valid values for this attribute are True and False. The default value is False (which means that by default a server instance boots in development mode.)		
	Note: If you boot the server in production mode by setting this attribute to True, you must reboot the server to set the mode back to development mode. Or in other words, you cannot reset the mode on a running server using other administrative tools, such as the WebLogic Server Scripting Tool (WLST).		
host	The DNS name or IP address on which the server instance is listening.	String	No
	The default value for this attribute is localhost.		
port	The TCP port number on which the server instance is listening.	int	No
	The default value for this attribute is 7001.		
generateconfig	Specifies whether or not wlserver creates a new domain for the specified server.	boolean	No
	Valid values for this attribute are true and false. The default value is false.		

Table 2-1 Attributes of the wlserver Ant Task (Continued)

Attribute	Description	Data Type	Required?
action	Specifies the action wlserver performs: start, shutdown, reboot, or connect. The shutdown action can be used with the optional forceshutdown attribute perform a forced shutdown. The default value for this attribute is start.	String	No
failonerror	This is a global attribute used by WebLogic Server Ant tasks. It specifies whether the task should fail if it encounters an error during the build. Valid values for this attribute are true and false. The default value is false.	Boolean	No

forceshutdown	This optional attribute is used in conjunction with the action="shutdown" attribute to perform a forced shutdown. For example:
	
	<pre><wlserver
	  host="${wls.host}"
	  port="${port}"
	  username="${wls.username}"
	  password="${wls.password}"
	  action="shutdown"
	  forceshutdown="true"/></pre>Valid values for this attribute are true and false. The default value is false.
		Boolean	No

noExit	(Optional) Leave the server process running after ant exits. Valid values are true or false. The default value is false, which means the server process will shut down when ant exits.	Boolean	No
protocol	Specifies the protocol that the wlserver Ant task uses to communicate with the WebLogic Server instance. Valid values are t3, t3s, http, https, and iiop. The default value is t3.	String	No

Table 2-1 Attributes of the wlserver Ant Task (Continued)

Attribute	Description	Data Type	Required?
forceImplicitUpgrade	Specifies whether the wlserver Ant task, if run against an 8.1 (or previous) domain, should implicitly upgrade it. Valid values are true or false. The default value is false, which means that the Ant task does *not* implicitly upgrade the domain, but rather, will fail with an error indicating that the domain needs to be upgraded. For more information about upgrading domains, see Upgrading WebLogic Application Environments.	Boolean	No.
configFile	Specifies the configuration file for your domain. The value of this attribute must be a valid XML file that conforms to the XML schema as defined in the WebLogic Server Domain Configuration Schema Reference. The XML file must exist in the Administration Server's root directory, which is either the current directory or the directory that you specify with the dir attribute. If you do not specify this attribute, the default value is config.xml in the directory specified by the dir attribute. If you do not specify the dir attribute, then the default domain directory is the current directory.	String	No.

Table 2-1 Attributes of the wlserver Ant Task (Continued)

Attribute	Description	Data Type	Required?
useBootProperties	Specifies whether to use the `boot.properties` file when starting a WebLogic Server instance. If this attribute is set to `true`, WebLogic Server uses the username and encrypted password stored in the `boot.properties` file to start rather than any values set with the `username` and `password` attributes. **Note:** The values of the `username` and `password` attributes are still used when shutting down or rebooting the WebLogic Server instance. The `useBootProperties` attribute applies *only* when starting the server. Valid values for this attribute are `true` and `false`. The default value is `false`.	Boolean	No
verbose	Specifies that the Ant task output additional information as it is performing its action. Valid values for this attribute are `true` and `false`. The default value is `false`.	Boolean	No

Configuring a WebLogic Server Domain Using the wlconfig Ant Task

The following sections describe how to use the `wlconfig` Ant task to configure a WebLogic Server domain.

WARNING: The `wlconfig` Ant task works *only* against MBeans that are compatible with the MBean server, which was deprecated as of version 9.0 of WebLogic Server. In particular, the `wlconfig` Ant task uses the deprecated proprietary API `weblogic.management.MBeanHome` to access WebLogic MBeans; therefore, `wlconfig` does *not* use the standard JMX interface (`javax.management.MBeanServerConnection`) to discover MBeans. This means that the only MBeans that you can access using `wlconfig` are those listed under the Deprecated MBeans category in the WebLogic Server MBean Reference.

For equivalent functionality, you should use the WebLogic Scripting Tool (WLST). See the WebLogic Scripting Tool documentation.

What the wlconfig Ant Task Does

The `wlconfig` Ant task enables you to configure a WebLogic Server domain by creating, querying, or modifying configuration MBeans on a running Administration Server instance. Specifically, `wlconfig` enables you to:

- Create new MBeans, optionally storing the new MBean Object Names in Ant properties.

- Set attribute values on a named MBean available on the Administration Server.

- Create MBeans and set their attributes in one step by nesting set attribute commands within create MBean commands.

- Query MBeans, optionally storing the query results in an Ant property reference.

- Query MBeans and set attribute values on all matching results.

- Establish a parent/child relationship among MBeans by nesting create commands within other create commands.

Basic Steps for Using wlconfig

1. Set your environment in a command shell. See "Basic Steps for Using wlserver" on page 2-2 for details.

   Note: The `wlconfig` task is predefined in the version of Ant shipped with WebLogic Server. If you want to use the task with your own Ant installation, add the following task definition in your build file:

   ```
   <taskdef name="wlconfig"
   classname="weblogic.ant.taskdefs.management.WLConfig"/>
   ```

2. `wlconfig` is commonly used in combination with `wlserver` to configure a new WebLogic Server domain created in the context of an Ant task. If you will be using `wlconfig` to configure such a domain, first use `wlserver` attributes to create a new domain and start the WebLogic Server instance.

3. Add an initial call to the `wlconfig` task to connect to the Administration Server for a domain. For example:

   ```
   <target name="doconfig">
      <wlconfig url="t3://localhost:7001" username="weblogic"
   ```

```
        password="weblogic">
</target>
```

4. Add nested `create`, `delete`, `get`, `set`, and `query` elements to configure the domain.

5. Execute the Ant task or tasks specified in the `build.xml` file by typing `ant` in the staging directory, optionally passing the command a target argument:

```
prompt> ant doconfig
```

Use `ant -verbose` to obtain more detailed messages from the `wlconfig` task.

Tip: Since WLST is the recommended tool for domain creation scripts, you should refer to the WLST offline sample scripts that are installed with the software. The offline scripts demonstrate how to create domains using the domain templates and are located in the following directory: *WL_HOME*\common\templates\scripts\wlst, where *WL_HOME* refers to the top-level installation directory for WebLogic Server. For example, the `basicWLSDomain.py` script creates a simple WebLogic domain, while `sampleMedRecDomain.py` creates a domain that defines resources similar to those used in the Avitek MedRec sample. See the WebLogic Scripting Tool documentation.

---

wlconfig Ant Task Reference

The following sections describe the attributes and elements that can be used with `wlconfig`.

Main Attributes

The following table describes the main attributes of the `wlconfig` Ant task.

Table 2-2 Main Attributes of the wlconfig Ant Task

Attribute	Description	Data Type	Required?
url	The URL of the domain's Administration Server.	String	Yes
username	The username of an administrator account.	String	No

Table 2-2 Main Attributes of the wlconfig Ant Task

Attribute	Description	Data Type	Required?
password	The password of an administrator account.	String	No
	To avoid having the plain text password appear in the build file or in process utilities such as ps, first store a valid username and encrypted password in a configuration file using the WebLogic Scripting Tool (WLST) storeUserConfig command. Then omit both the username and password attributes in your Ant build file. When the attributes are omitted, wlconfig attempts to login using values obtained from the default configuration file.		
	If you want to obtain a username and password from a non-default configuration file and key file, use the userconfigfile and userkeyfile attributes with wlconfig.		
	See the command reference for storeUserConfig in the WLST Command and Variable Reference for more information on storing and encrypting passwords.		
failonerror	This is a global attribute used by WebLogic Server Ant tasks. It specifies whether the task should fail if it encounters an error during the build. This attribute is set to true by default.	Boolean	No

Table 2-2 Main Attributes of the wlconfig Ant Task

Attribute	Description	Data Type	Required?
userconfigfile	Specifies the location of a user configuration file to use for obtaining the administrative username and password. Use this option, instead of the username and password attributes, in your build file when you do not want to have the plain text password shown in-line or in process-level utilities such as ps. Before specifying the userconfigfile attribute, you must first generate the file using the WebLogic Scripting Tool (WLST) storeUserConfig command as described in the WLST Command and Variable Reference.	File	No
userkeyfile	Specifies the location of a user key file to use for encrypting and decrypting the username and password information stored in a user configuration file (the userconfigfile attribute). Before specifying the userkeyfile attribute, you must first generate the key file using the WebLogic Scripting Tool (WLST) storeUserConfig command as described in the WLST Command and Variable Reference.	File	No

Nested Elements

wlconfig also has several elements that can be nested to specify configuration options:

- create

- delete

- set

- get

- query

- invoke

create

The `create` element creates a new MBean in the WebLogic Server domain. The `wlconfig` task can have any number of `create` elements.

A `create` element can have any number of nested `set` elements, which set attributes on the newly-created MBean. A `create` element may also have additional, nested `create` elements that create child MBeans.

The `create` element has the following attributes.

Table 2-3 Attributes of the create Element

Attribute	Description	Data Type	Required?
name	The name of the new MBean object to create.	String	No (`wlconfig` supplies a default name if none is specified.)
type	The MBean type.	String	Yes
property	The name of an optional Ant property that holds the object name of the newly-created MBean. **Note:** If you nest a `create` element inside of another `create` element, you cannot specify the `property` attribute for the *nested* `create` element.	String	No

delete

The delete element removes an existing MBean from the WebLogic Server domain. delete takes a single attribute:

Table 2-4 Attribute of the delete Element

Attribute	Description	Data Type	Required?
mbean	The object name of the MBean to delete.	String	Required when the delete element is a direct child of the wlconfig task. Not required when nested within a query element.

set

The set element sets MBean attributes on a named MBean, a newly-created MBean, or on MBeans retrieved as part of a query. You can include the set element as a direct child of the wlconfig task, or nested within a create or query element.

The set element has the following attributes:

Table 2-5 Attributes of the set Element

Attribute	Description	Data Type	Required?
attribute	The name of the MBean attribute to set.	String	Yes
value	The value to set for the specified MBean attribute. You can specify multiple object names (stored in Ant properties) as a value by delimiting the entire value list with quotes and separating the object names with a semicolon.	String	Yes

Table 2-5 Attributes of the set Element

Attribute	Description	Data Type	Required?
mbean	The object name of the MBean whose values are being set. This attribute is required only when the `set` element is included as a direct child of the main `wlconfig` task; it is not required when the set element is nested within the context of a `create` or `query` element.	String	Required only when the `set` element is a direct child of the `wlconfig` task.
domain	This attribute specifies the JMX domain name for Security MBeans and third-party SPI MBeans. It is not required for administration MBeans, as the domain corresponds to the WebLogic Server domain. **Note:** You cannot use this attribute if the `set` element is nested inside of a `create` element.	String	No

get

The `get` element retrieves attribute values from an MBean in the WebLogic Server domain. The `wlconfig` task can have any number of `get` elements.

The `get` element has the following attributes.

Table 2-6 Attributes of the get Element

Attribute	Description	Data Type	Required?
attribute	The name of the MBean attribute whose value you want to retrieve.	String	Yes

Table 2-6 Attributes of the get Element

Attribute	Description	Data Type	Required?
property	The name of an Ant property that will hold the retrieved MBean attribute value.	String	Yes
mbean	The object name of the MBean you want to retrieve attribute values from.	String	Yes

query

The `query` elements finds MBean that match a search pattern.

The query element supports the following nested child elements:

- `set`—performs set operations on all MBeans in the result set.

- `get`—performs get operations on all MBeans in the result set.

- `create`—each MBean in the result set is used as a parent of a new MBean.

- `delete`—performs delete operations on all MBeans in the result set.

- `invoke`—invokes all matching MBeans in the result set.

`wlconfig` can have any number of nested `query` elements.

`query` has the following attributes:

Table 2-7 Attributes of the query Element

Attribute	Description	Data Type	Required?
domain	The name of the WebLogic Server domain in which to search for MBeans.	String	No
type	The type of MBean to query.	String	No
name	The name of the MBean to query.	String	No
pattern	A JMX query pattern.	String	No

Table 2-7 Attributes of the query Element

Attribute	Description	Data Type	Required?
property	The name of an optional Ant property that will store the query results.	String	No
domain	This attribute specifies the JMX domain name for Security MBeans and third-party SPI MBeans. It is not required for administration MBeans, as the domain corresponds to the WebLogic Server domain.	String	No

invoke

The `invoke` element invokes a management operation for one or more MBeans. For WebLogic Server MBeans, you usually use this command to invoke operations other than the `getAttribute` and `setAttribute` that most WebLogic Server MBeans provide.

The `invoke` element has the following attributes.

Table 2-8 Attributes of the invoke Element

Attribute	Description	Data Type	Required?
mbean	The object name of the MBean you want to invoke.	String	You must specify either the `mbean` or `type` attribute of the invoke element.
type	The type of MBean to invoke.	String	You must specify either the `mbean` or `type` attribute of the invoke element.

Table 2-8 Attributes of the invoke Element

Attribute	Description	Data Type	Required?
methodName	The method of the MBean to invoke.	String	Yes
arguments	The list of arguments (separated by spaces) to pass to the method specified by the methodName attribute.	String	No

Using the libclasspath Ant Task

Use the libclasspath Ant task to build applications that use libraries, such as application libraries and web libraries.

-
-
-

libclasspath Task Definition

To use the task with your own Ant installation, add the following task definition in your build file:

```
<taskdef name="libclasspath" classname="weblogic.ant.taskdefs.build.Lib
ClasspathTask"/>
```

libclasspath Ant Task Reference

The following sections describe the attributes and elements that can be used with the libclasspath Ant task.

-
-

Main libclasspath Attributes

The following table describes the main attributes of the libclasspath Ant task.

Table 2-9 Attributes of the libclasspath Ant Task

Attribute	Description	Required
basedir	The root of .ear or .war file to extract from.	One of the two attributes is required.
basewar	The name of the .war file to extract from.	If basewar is specified, basedir is ignored and the library referenced in basewar is used as the .war file to extract classpath or resourcepath information from.
tmpdir	The fully qualified name of the directory to be used for extracting libraries.	Yes.
classpathproperty	Contains the classpath for the referenced libraries. For example, if basedir points to a .war file that references web application libraries in the weblogic.xml file, the classpathproperty contains the WEB-INF/classes and WEB-INF/lib directories of the web application libraries. Additionally, if basedir points to a .war file that has .war files under WEB-INF/bea-ext, the classpathproperty contains the WEB-INF/classes and WEB-INF/lib directories for the Oracle extensions.	At least one of the two attributes is required.
resourcepathproperty	Contains library resources that are not classes. For example, if basedir points to a .war file that has .war files under WEB-INF/bea-ext, resourcepathproperty contains the roots of the exploded extensions.	

Nested libclasspath Elements

libclasspath also has two elements that can be nested to specify configuration options. At least one of the elements is required when using the libclasspath Ant task:

librarydir

The following attribute is required when using this element:

 dir—Specifies that all files in this directory are registered as available libraries.

library

The following attribute is required when using this element:

 file—Register this file as an available library.

Example libclasspath Ant Task

This section provides example code of a libclasspath Ant task:

Listing 2-1   Example libclasspath Ant Task Code

```
.
.
.
   <taskdef name="libclasspath" classname="weblogic.ant.taskdefs.build.Lib
ClasspathTask"/>

   <!-- Builds classpath based on libraries defined in weblogic-applicatio
n.xml. -->
   <target name="init.app.libs">
      <libclasspath basedir="${src.dir}" tmpdir="${tmp.dir}" classpathprop
erty="app.lib.classpath">
         <librarydir dir="${weblogic.home}/common/deployable-libraries/"/>
      </libclasspath>
   <echo message="app.lib.claspath is ${app.lib.classpath}" level="info"/>
   </target>
```

-
-
-

Creating a Split Development Directory Environment

The following sections describe the steps for creating a WebLogic Server split development directory that you can use to develop a Java EE application or module:

Overview of the Split Development Directory Environment

The WebLogic split development directory environment consists of a directory layout and associated Ant tasks that help you repeatedly build, change, and deploy Java EE applications. Compared to other development frameworks, the WebLogic split development directory provides these benefits:

- **Fast development and deployment.** By minimizing unnecessary file copying, the split development directory Ant tasks help you recompile and redeploy applications quickly *without* first generating a deployable archive file or exploded archive directory.

- **Simplified build scripts.** The Oracle-provided Ant tasks automatically determine which Java EE modules and classes you are creating, and build components in the correct order to support common classpath dependencies. In many cases, your project build script can simply identify the source and build directories and allow Ant tasks to perform their default behaviors.

- **Easy integration with source control systems.** The split development directory provides a clean separation between source files and generated files. This helps you maintain only editable files in your source control system. You can also clean the build by deleting the entire build directory; build files are easily replaced by rebuilding the project.

Source and Build Directories

The source and build directories form the basis of the split development directory environment. The source directory contains all editable files for your project—Java source files, editable descriptor files, JSPs, static content, and so forth. You create the source directory for an application by following the directory structure guidelines described in "Organizing Java EE Components in a Split Development Directory" on page 3-6.

The top level of the source directory always represents an Enterprise Application (`.ear` file), even if you are developing only a single Java EE module. Subdirectories beneath the top level source directory contain:

- Enterprise Application Modules (EJBs and Web Applications)

  Note: The split development directory structure does not provide support for developing new Resource Adapter components.

- Descriptor files for the Enterprise Application (`application.xml` and `weblogic-application.xml`)

- Utility classes shared by modules of the application (for example, exceptions, constants)

- Libraries (compiled `.jar` files, including third-party libraries) used by modules of the application

The build directory contents are generated automatically when you run the `wlcompile` ant task against a valid source directory. The `wlcompile` task recognizes EJB, Web Application, and shared library and class directories in the source directory, and builds those components in an

order that supports common class path requirements. Additional Ant tasks can be used to build Web Services or generate deployment descriptor files from annotated EJB code.

Figure 3-1 Source and Build Directories
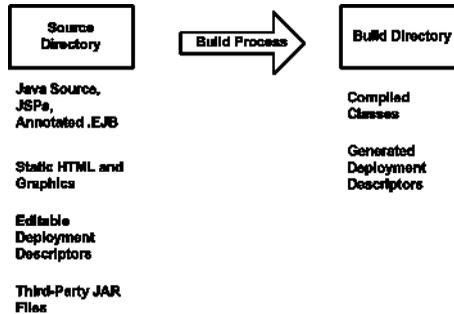
The build directory contains only those files generated during the build process. The combination of files in the source and build directories form a deployable Java EE application.

The build and source directory contents can be place in any directory of your choice. However, for ease of use, the directories are commonly placed in directories named `source` and `build`, within a single project directory (for example, `\myproject\build` and `\myproject\source`).

Deploying from a Split Development Directory

All WebLogic Server deployment tools (`weblogic.Deployer`, `wldeploy`, and the Administration Console) support direct deployment from a split development directory. You specify only the build directory when deploying the application to WebLogic Server.

WebLogic Server attempts to use all classes and resources available in the *source* directory for deploying the application. If a required resource is not available in the source directory, WebLogic Server then looks in the application's build directory for that resource. For example, if a deployment descriptor is generated during the build process, rather than stored with source code as an editable file, WebLogic Server obtains the generated file from the build directory.

WebLogic Server discovers the location of the source directory by examining the `.beabuild.txt` file that resides in the top level of the application's build directory. If you ever move or modify the source directory location, edit the `.beabuild.txt` file to identify the new source directory name.

"Deploying and Packaging from a Split Development Directory" on page 5-1 describes the `wldeploy` Ant task that you can use to automate deployment from the split directory environment.

Figure 3-2 shows a typical deployment process. The process is initiated by specifying the build directory with a WebLogic Server tool. In the figure, all compiled classes and generated deployment descriptors are discovered in the build directory, but other application resources (such as static files and editable deployment descriptors) are missing. WebLogic Server uses the hidden `.beabuild.txt` file to locate the application's source directory, where it finds the required resources.

Figure 3-2   Split Directory Deployment
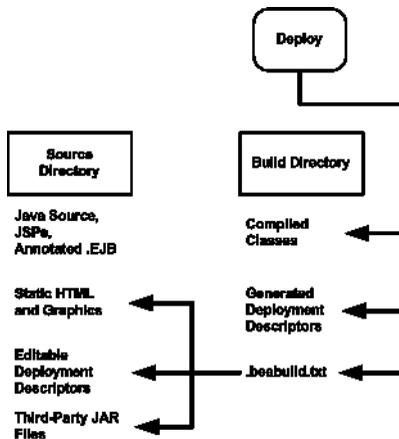
Split Development Directory Ant Tasks

Oracle provides a collection of Ant tasks designed to help you develop applications using the split development directory environment. Each Ant task uses the source, build, or both directories to perform common development tasks:

- `wlcompile`—This Ant task compiles the contents of the source directory into subdirectories of the build directory. `wlcompile` compiles Java classes and also processes

annotated `.ejb` files into deployment descriptors, as described in "Compiling Applications Using wlcompile" on page 4-1.

- `wlappc`—This Ant task invokes the appc compiler, which generates JSPs and container-specific EJB classes for deployment. See "Building Modules and Applications Using wlappc" on page 4-4.

- `wldeploy`—This Ant task deploys any format of Java EE applications (exploded or archived) to WebLogic Server. To deploy directly from the split development directory environment, you specify the build directory of your application. See "wldeploy Ant Task Reference" on page B-1.

- `wlpackage`—This Ant task uses the contents of both the source and build directories to generate an EAR file or exploded EAR directory that you can give to others for deployment.

Using the Split Development Directory Structure: Main Steps

The following steps illustrate how you use the split development directory structure to build and deploy a WebLogic Server application.

1. Create the main EAR source directory for your project. When using the split development directory environment, you must develop Web Applications and EJBs as part of an Enterprise Application, even if you do not intend to develop multiple Java EE modules. See "Organizing Java EE Components in a Split Development Directory" on page 3-6.

2. Add one or more subdirectories to the EAR directory for storing the source for Web Applications, EJB components, or shared utility classes. See "Organizing Java EE Components in a Split Development Directory" on page 3-6 and "Organizing Shared Classes in a Split Development Directory" on page 3-11.

3. Store all of your editable files (source code, static content, editable deployment descriptors) for modules in subdirectories of the EAR directory. Add the entire contents of the source directory to your source control system, if applicable.

4. Set your WebLogic Server environment by executing either the `setWLSEnv.cmd` (Windows) or `setWLSEnv.sh` (UNIX) script. The scripts are located in the `WL_HOME\server\bin\` directory, where `WL_HOME` is the top-level directory in which WebLogic Server is installed.

5. Use the `weblogic.BuildXMLGen` utility to generate a default `build.xml` file for use with your project. Edit the default property values as needed for your environment. See "Generating a Basic build.xml File Using weblogic.BuildXMLGen" on page 3-13.

6. Use the default targets in the `build.xml` file to build, deploy, and package your application. See "Generating a Basic build.xml File Using weblogic.BuildXMLGen" on page 3-13 for a list of default targets.

Organizing Java EE Components in a Split Development Directory

The split development directory structure requires each project to be staged as a Java EE Enterprise Application. Oracle therefore recommends that you stage even stand-alone Web applications and EJBs as modules of an Enterprise application, to benefit from the split directory Ant tasks. This practice also allows you to easily add or remove modules at a later date, because the application is already organized as an EAR.

Note: If your project requires multiple EARs, see also "Developing Multiple-EAR Projects Using the Split Development Directory" on page 3-15.

The following sections describe the basic conventions for staging the following module types in the split development directory structure:

- "Enterprise Application Configuration" on page 3-8

- "Web Applications" on page 3-8

- "EJBs" on page 3-10

- "Shared Utility Classes" on page 3-11

- "Third-Party Libraries" on page 3-12

The directory examples are taken from the `splitdir` sample application installed in `WL_HOME\samples\server\examples\src\examples\splitdir`, where `WL_HOME` is your WebLogic Server installation directory.

Source Directory Overview

The following figure summarizes the source directory contents of an Enterprise Application having a Web Application, EJB, shared utility classes, and third-party libraries. The sections that follow provide more details about how individual parts of the enterprise source directory are organized.

Figure 3-3   Overview of Enterprise Application Source Directory
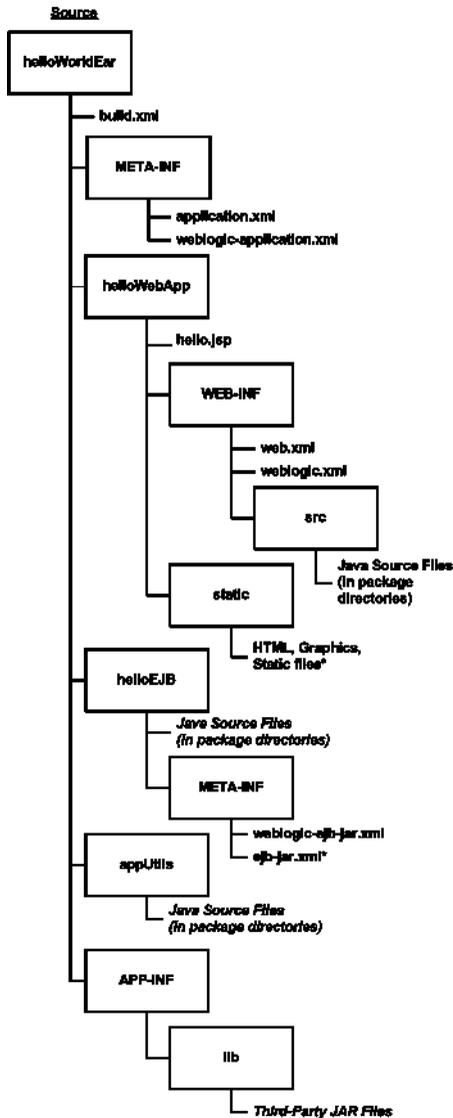
Enterprise Application Configuration

The top level source directory for a split development directory project represents an Enterprise Application. The following figure shows the minimal files and directories required in this directory.

Figure 3-4   Enterprise Application Source Directory
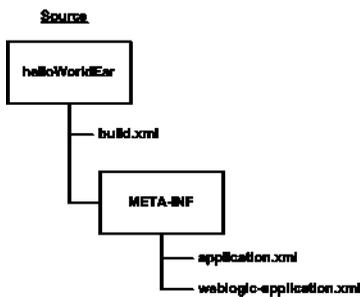
The Enterprise Application directory will also have one or more subdirectories to hold a Web Application, EJB, utility class, and/or third-party Jar file, as described in the following sections.

Web Applications

Web Applications use the basic source directory layout shown in the figure below.

Figure 3-5 Web Application Source and Build Directories
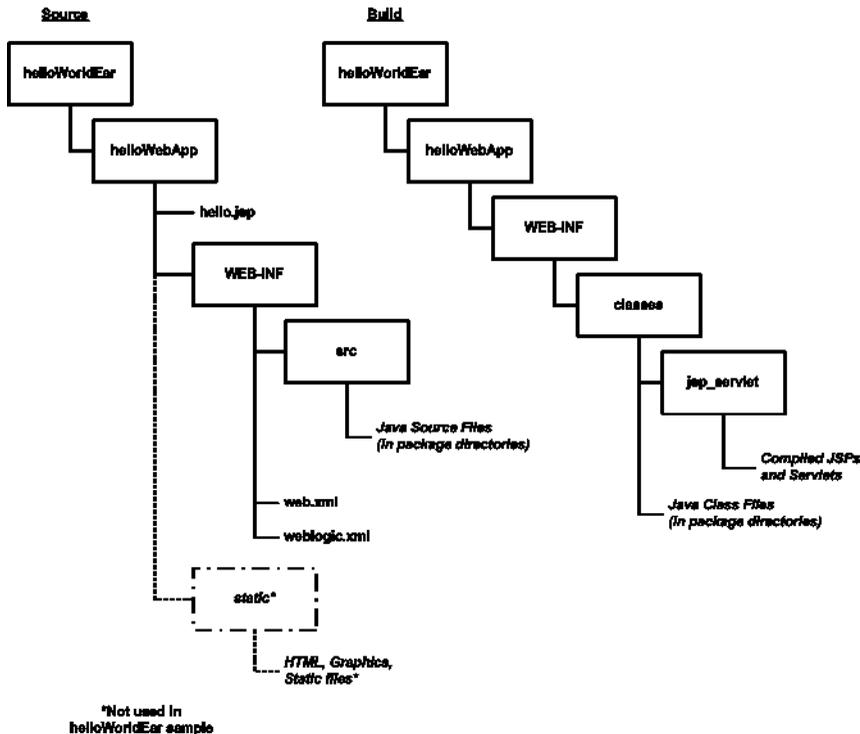
The key directories and files for the Web Application are:

- `helloWebApp\` —The top level of the Web Application module can contain JSP files and static content such as HTML files and graphics used in the application. You can also store static files in any named subdirectory of the Web Application (for example, `helloWebApp\graphics` or `helloWebApp\static`.)

- `helloWebApp\WEB-INF\` —Store the Web Application's editable deployment descriptor files (`web.xml` and `weblogic.xml`) in the `WEB-INF` subdirectory.

- `helloWebApp\WEB-INF\src` —Store Java source files for Servlets in package subdirectories under `WEB-INF\src`.

When you build a Web Application, the `appc` Ant task and `jspc` compiler compile JSPs into package subdirectories under `helloWebApp\WEB-INF\classes\jsp_servlet` in the build directory. Editable deployment descriptors are not copied during the build process.

EJBs

EJBs use the source directory layout shown in the figure below.

Figure 3-6   EJB Source and Build Directories
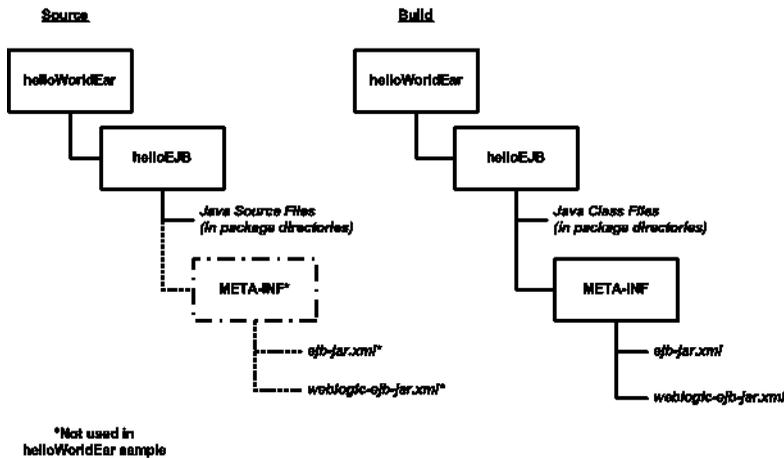
The key directories and files for an EJB are:

- `helloEJB\` —Store all EJB source files under package directories of the EJB module directory. The source files can be either `.java` source files, or annotated `.ejb` files.

- `helloEJB\META-INF\` —Store editable EJB deployment descriptors (ejb-jar.xml and weblogic-ejb-jar.xml) in the META-INF subdirectory of the EJB module directory. The `helloWorldEar` sample does not include a `helloEJB\META-INF` subdirectory, because its deployment descriptors files are generated from annotations in the `.ejb` source files. See "Important Notes Regarding EJB Descriptors" on page 3-11.

During the build process, EJB classes are compiled into package subdirectories of the `helloEJB` module in the build directory. If you use annotated `.ejb` source files, the build process also generates the EJB deployment descriptors and stores them in the `helloEJB\META-INF` subdirectory of the build directory.

Important Notes Regarding EJB Descriptors

EJB deployment descriptors should be included in the source META-INF directory and treated as source code *only* if those descriptor files are created from scratch or are edited manually. Descriptor files that are generated from annotated .ejb files should appear only in the build directory, and they can be deleted and regenerated by building the application.

For a given EJB component, the EJB source directory should contain either:

- EJB source code in .java source files and editable deployment descriptors in META-INF

or:

- EJB source code with descriptor annotations in .ejb source files, and *no editable descriptors* in META-INF.

In other words, do not provide both annotated .ejb source files and editable descriptor files for the same EJB component.

Organizing Shared Classes in a Split Development Directory

The WebLogic split development directory also helps you store shared utility classes and libraries that are required by modules in your Enterprise Application. The following sections describe the directory layout and classloading behavior for shared utility classes and third-party JAR files.

Shared Utility Classes

Enterprise Applications frequently use Java utility classes that are shared among application modules. Java utility classes differ from third-party JARs in that the source files are part of the application and must be compiled. Java utility classes are typically libraries used by application modules such as EJBs or Web applications.

Figure 3-7   Java Utility Class Directory
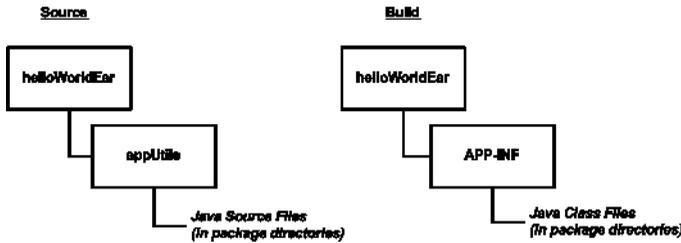
Place the source for Java utility classes in a named subdirectory of the top-level Enterprise Application directory. Beneath the named subdirectory, use standard package subdirectory conventions.

During the build process, the `wlcompile` Ant task invokes the `javac` compiler and compiles Java classes into the `APP-INF/classes/` directory under the build directory. This ensures that the classes are available to other modules in the deployed application.

Third-Party Libraries

You can extend an Enterprise Application to use third-party `.jar` files by placing the files in the `APP-INF\lib\` directory, as shown below:

Figure 3-8   Third-party Library Directory
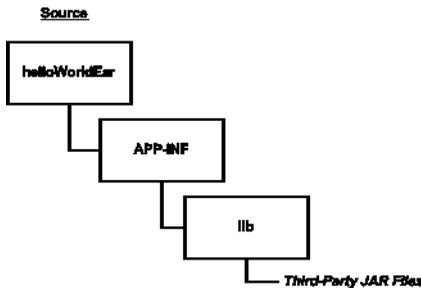
Third-party JARs are generally not compiled, but may be versioned using the source control system for your application code. For example, XML parsers, logging implementations, and Web Application framework JAR files are commonly used in applications and maintained along with editable source code.

During the build process, third-party JAR files are not copied to the build directory, but remain in the source directory for deployment.

Class Loading for Shared Classes

The classes and libraries stored under `APP-INF/classes` and `APP-INF/lib` are available to all modules in the Enterprise Application. The application classloader always attempts to resolve class requests by first looking in `APP-INF/classes`, then `APP-INF/lib`.

Generating a Basic build.xml File Using weblogic.BuildXMLGen

After you set up your source directory structure, use the `weblogic.BuildXMLGen` utility to create a basic `build.xml` file. `weblogic.BuildXMLGen` is a convenient utility that generates an Ant `build.xml` file for Enterprise applications that are organized in the split development directory structure. The utility analyzes the source directory and creates build and deploy targets for the Enterprise application as well as individual modules. It also creates targets to clean the build and generate new deployment descriptors.

Additionally, optional packages are supported as Java EE shared libraries in `weblogic.BuildXMLGen`, whereby all manifests of an application and its modules are scanned to look for optional package references. If optional package references are found they are added to the compile and `appc` tasks in the generated `build.xml` file.

For example, if a library located at `lib\echolib.jar` is referenced as an optional package, the tasks generated by `weblogic.BuildXMLGen` will contains an `appc` task that would appear as follows:

```
<target name="appc" description="Runs weblogic.appc on your application">
  <wlappc source="${dest.dir}" verbose="${verbose}">
    <library file="lib\echolib\echolib.jar" />
  </wlappc>
</target>
```

The compile and `appc` tasks for modules also use the `lib\echolib\echolib.jar` library.

weblogic.BuildXMLGen Syntax

The syntax for `weblogic.BuildXMLGen` is as follows:

```
java weblogic.BuildXMLGen [options] <source directory>
```

where `options` include:

- -help—print standard usage message

- -version—print version information

- `-projectName <project name>`—name of the Ant project

- `-d <directory>`—directory where `build.xml` is created. The default is the current directory.

- `-file <build.xml>`—name of the generated build file

- `-librarydir <directories>`—create build targets for shared Java EE libraries in the comma-separated list of directories. See "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.

- `-username <username>`—user name for deploy commands

- `-password <password>`—user password

After running `weblogic.BuildXMLGen`, edit the generated `build.xml` file to specify properties for your development environment. The list of properties you need to edit are shown in the listing below.

Listing 3-1   build.xml Editable Properties

```
<!-- BUILD PROPERTIES ADJUST THESE FOR YOUR ENVIRONMENT -->
  <property name="tmp.dir" value="/tmp" />
  <property name="dist.dir" value="${tmp.dir}/dist"/>
  <property name="app.name" value="helloWorldEar" />
  <property name="ear" value="${dist.dir}/${app.name}.ear"/>
  <property name="ear.exploded" value="${dist.dir}/${app.name}_exploded"/>
  <property name="verbose" value="true" />
  <property name="user" value="USERNAME" />
  <property name="password" value="PASSWORD" />
  <property name="servername" value="myserver" />
  <property name="adminurl" value="iiop://localhost:7001" />
```

In particular, make sure you edit the `tmp.dir` property to point to the build directory you want to use. By default, the `build.xml` file builds projects into a subdirectory `tmp.dir` named after the application (`/tmp/helloWorldEar` in the above listing).

The following listing shows the default main targets created in the `build.xml` file. You can view these targets at the command prompt by entering the `ant -projecthelp` command in the EAR source directory.

Listing 3-2   Default build.xml Targets

```
appc                 Runs weblogic.appc on your application
build                Compiles helloWorldEar application and runs appc
clean                Deletes the build and distribution directories
compile              Only compiles helloWorldEar application, no appc
compile.appStartup   Compiles just the appStartup module of the application
compile.appUtils     Compiles just the appUtils module of the application
compile.build.orig   Compiles just the build.orig module of the application
compile.helloEJB     Compiles just the helloEJB module of the application
compile.helloWebApp  Compiles just the helloWebApp module of the application
compile.javadoc      Compiles just the javadoc module of the application
deploy              Deploys (and redeploys) the entire helloWorldEar application
descriptors          Generates application and module descriptors
ear                  Package a standard J2EE EAR for distribution
ear.exploded         Package a standard exploded J2EE EAR
redeploy.appStartup  Redeploys just the appStartup module of the application
redeploy.appUtils    Redeploys just the appUtils module of the application
redeploy.build.orig  Redeploys just the build.orig module of the application
redeploy.helloEJB    Redeploys just the helloEJB module of the application
redeploy.helloWebApp Redeploys just the helloWebApp module of application
redeploy.javadoc     Redeploys just the javadoc module of the application
undeploy             Undeploys the entire helloWorldEar application
```

Developing Multiple-EAR Projects Using the Split Development Directory

The split development directory examples and procedures described previously have dealt with projects consisting of a single Enterprise Application. Projects that require building multiple Enterprise Applications simultaneously require slightly different conventions and procedures, as described in the following sections.

Note: The following sections refer to the MedRec sample application, which consists of three separate Enterprise Applications as well as shared utility classes, third-party JAR files, and dedicated client applications. The MedRec source and build directories are installed under *WL_HOME*/samples/server/medrec, where *WL_HOME* is the WebLogic Server installation directory.

Organizing Libraries and Classes Shared by Multiple EARs

For single EAR projects, the split development directory conventions suggest keeping third-party JAR files in the APP-INF/lib directory of the EAR source directory. However, a multiple-EAR project would require you to maintain a copy of the same third-party JAR files in the APP-INF/lib directory of *each* EAR source directory. This introduces multiple copies of the source JAR files, increases the possibility of some JAR files being at different versions, and requires additional space in your source control system.

To address these problems, consider editing your build script to copy third-party JAR files into the APP-INF/lib directory of the *build* directory for each EAR that requires the libraries. This allows you to maintain a single copy and version of the JAR files in your source control system, yet it enables each EAR in your project to use the JAR files.

The MedRec sample application installed with WebLogic Server uses this strategy, as shown in the following figure.

Figure 3-9   Shared JAR Files in MedRec
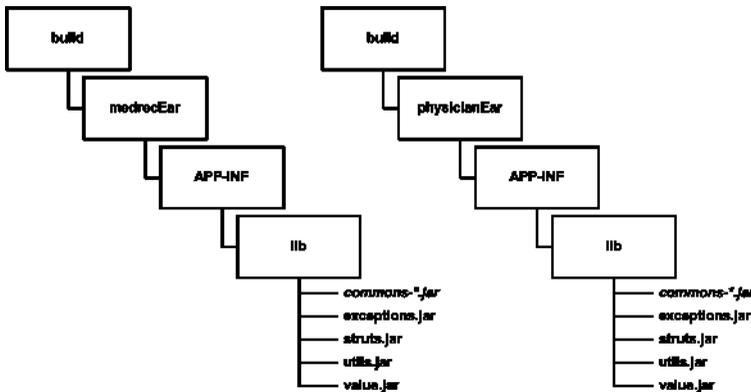
MedRec takes a similar approach to utility classes that are shared by multiple EARs in the project. Instead of including the source for utility classes within the scope of each ear that needs them, MedRec keeps the utility class source independent of all EARs. After compiling the utility classes, the build script archives them and copies the JARs into the build directory under the APP-INF/LIB subdirectory of each EAR that uses the classes, as shown in figure Figure 3-9.

Linking Multiple build.xml Files

When developing multiple EARs using the split development directory, each EAR project generally uses its own `build.xml` file (perhaps generated by multiple runs of `weblogic.BuildXMLGen`.). Applications like MedRec also use a master `build.xml` file that calls the subordinate `build.xml` files for each EAR in the application suite.

Ant provides a core task (named `ant`) that allows you to execute other project build files within a master `build.xml` file. The following line from the MedRec master build file shows its usage:

```
<ant inheritAll="false" dir="${root}/startupEar" antfile="build.xml"/>
```

The above task instructs Ant to execute the file named `build.xml` in the `/startupEar` subdirectory. The `inheritAll` parameter instructs Ant to pass only user properties from the master build file tot the `build.xml` file in `/startupEar`.

MedRec uses multiple tasks similar to the above to build the `startupEar`, `medrecEar`, and `physicianEar` applications, as well as building common utility classes and client applications.

Best Practices for Developing WebLogic Server Applications

Oracle recommends the following "best practices" for application development.

- Package applications as part of an Enterprise application. See "Packaging Applications Using wlpackage" on page 5-1.

- Use the split development directory structure. See "Organizing Java EE Components in a Split Development Directory" on page 3-6.

- For distribution purposes, package and deploy in archived format. See "Packaging Applications Using wlpackage" on page 5-1.

- In most other cases, it is more convenient to deploy in exploded format. See "Archive versus Exploded Archive Directory" on page 5-2.

- Never deploy untested code on a WebLogic Server instance that is serving production applications. Instead, set up a development WebLogic Server instance on the same computer on which you edit and compile, or designate a WebLogic Server development location elsewhere on the network.

- Even if you do not run a development WebLogic Server instance on your development computer, you must have access to a WebLogic Server distribution to compile your

programs. To compile any code using WebLogic or Java EE APIs, the Java compiler needs access to the `weblogic.jar` file and other JAR files in the distribution directory. Install WebLogic Server on your development computer to make WebLogic distribution files available locally.

Building Applications in a Split Development Directory

The following sections describe the steps for building WebLogic Server Java EE applications using the WebLogic split development directory environment:

- "Compiling Applications Using wlcompile" on page 4-1
- "Building Modules and Applications Using wlappc" on page 4-4

Compiling Applications Using wlcompile

You use the `wlcompile` Ant task to invoke the `javac` compiler to compile your application's Java components in a split development directory structure. The basic syntax of `wlcompile` identifies the source and build directories, as in this command from the `helloWorldEar` sample:

```
<wlcompile srcdir="${src.dir}" destdir="${dest.dir}"/>
```

Note: Deployment descriptors are no longer mandatory in JEE 5; therefore, exploded module directories must indicate the module type by using the .war or .jar suffix when there is no deployment descriptor in these directories. The suffix is required so that `wlcompile` can recognize the modules. The .war suffix indicates the module is a Web application module and the .jar suffix indicates the module is an EJB module.

The following is the order in which events occur using this task:

1. `wlcompile` compiles the Java components into an output directory:

   `WL_HOME\samples\server\examples\build\helloWorldEar\APP-INF\classes\`

   where *WL_HOME* is the WebLogic Server installation directory.

2. `wlcompile` builds the EJBs and automatically includes the previously built Java modules in the compiler's classpath. This allows the EJBs to call the Java modules without requiring you to manually edit their classpath.

3. Finally, `wlcompile` compiles the Java components in the Web application with the EJB and Java modules in the compiler's classpath. This allows the Web applications to refer to the EJB and application Java classes without requiring you to manually edit the classpath.

Using includes and excludes Properties

More complex Enterprise applications may have compilation dependencies that are not automatically handled by the `wlcompile` task. However, you can use the include and exclude options to `wlcompile` to enforce your own dependencies. The `includes` and `excludes` properties accept the names of Enterprise Application modules—the names of subdirectories in the Enterprise application source directory—to include or exclude them from the compile stage.

The following line from the `helloWorldEar` sample shows the `appStartup` module being excluded from compilation:

```
<wlcompile srcdir="${src.dir}" destdir="${dest.dir}"
   excludes="appStartup"/>
```

wlcompile Ant Task Attributes

Table 4-1 contains Ant task attributes specific to `wlcompile`.

Table 4-1 wlcompile Ant Task Attributes

Attribute	Description
srcdir	The source directory.
destdir	The build/output directory.
classpath	Allows you to change the classpath used by `wlcompile`.
includes	Allows you to include specific directories from the build.
excludes	Allows you to exclude specific directories from the build.
librarydir	Specifies a directory of shared Java EE libraries to add to the classpath. See "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.

Nested javac Options

The `wlcompile` Ant task can accept nested `javac` options to change the compile-time behavior. For example, the following `wlcompile` command ignores deprecation warnings and enables debugging:

```
<wlcompile srcdir="${mysrcdir}" destdir="${mybuilddir}">
   <javac deprecation="false" debug="true"
   debuglevel="lines,vars,source"/>
</wlcompile>
```

Setting the Classpath for Compiling Code

Most WebLogic services are based on Java EE standards and are accessed through standard Java EE packages. The Sun, WebLogic, and other Java classes required to compile programs that use WebLogic services are packaged in the `weblogic.jar` file in the `lib` directory of your WebLogic Server installation. In addition to `weblogic.jar`, include the following in your compiler's CLASSPATH:

- The `lib\tools.jar` file in the JDK directory, or other standard Java classes required by the Java Development Kit you use.

- The `examples.property` file for Apache Ant (for examples environment). This file is discussed in the WebLogic Server documentation on building examples using Ant located at: `samples\server\examples\src\examples\examples.html`

- Classes for third-party Java tools or services your programs import.

- Other application classes referenced by the programs you are compiling.

Library Element for wlcompile and wlappc

The `library` element is an optional element used to define the name and optional version information for a module that represents a shared Java EE library required for building an application, as described in "Creating Shared Java EE Libraries and Optional Packages" on page 9-1. The `library` element can be used with both `wlcompile` and `wlappc`, described in "Building Modules and Applications Using wlappc" on page 4-4.

The name and version information are specified as attributes to the library element, described in Table 4-2.

Table 4-2 Library attributes

Attribute	Description
`file`	Required filename of a Java EE library
`name`	The optional name of a required Java EE library.
`specificationversion`	An optional specification version required for the library.
`implementationversion`	An optional implementation version required for the library.

The format choices for both `specificationversion` and `implementationversion` are described in "Referencing Shared Java EE Libraries in an Enterprise Application" on page 9-12. The following output shows a sample `library` reference:

```
<library file="c:\mylibs\lib.jar" name="ReqLib" specificationversion="920"
implementationversion="1.1" />
```

Building Modules and Applications Using wlappc

The `weblogic.appc` compiler generates JSPs and container-specific EJB classes for deployment, and validates deployment descriptors for compliance with the current Java EE specifications. `appc` performs validation checks between the application-level deployment descriptors and the individual modules in the application as well as validation checks across the modules.

Additionally, optional packages are supported as Java EE shared libraries in `appc`, whereby all manifests of an application and its modules are scanned to look for optional package references.

`wlappc` is the Ant task interface to the `weblogic.appc` compiler. The following section describe the `wlappc` options and usage. Both `weblogic.appc` and the `wlappc` Ant task compile modules in the order in which they appear in the `application.xml` deployment descriptor file that describes your Enterprise application.

wlappc Ant Task Attributes

Table 4-3 describes Ant task options specific to wlappc. These options are similar to the weblogic.appc command-line options, but with a few differences.

Notes: See "weblogic.appc Reference" on page 4-7 for a list of weblogic.appc options.

See also "Library Element for wlcompile and wlappc" on page 4-3.

Table 4-3 wlappc Ant Task Attributes

Option	Description
print	Prints the standard usage message.
version	Prints appc version information.
output <file>	Specifies an alternate output archive or directory. If not set, the output is placed in the source archive or directory.
forceGeneration	Forces generation of EJB and JSP classes. Without this flag, the classes may not be regenerated (if determined to be unnecessary).
lineNumbers	Adds line numbers to generated class files to aid in debugging.
writeInferredDescriptors	Specifies that the application or module contains deployment descriptors with annotation information.
basicClientJar	Does not include deployment descriptors in client JARs generated for EJBs.
idl	Generates IDL for EJB remote interfaces.
idlOverwrite	Always overwrites existing IDL files.
idlVerbose	Displays verbose information for IDL generation.
idlNoValueTypes	Does not generate valuetypes and the methods/attributes that contain them.
idlNoAbstractInterfaces	Does not generate abstract interfaces and methods/attributes that contain them.
idlFactories	Generates factory methods for valuetypes.

Table 4-3 wlappc Ant Task Attributes

`idlVisibroker`	Generates IDL somewhat compatible with Visibroker 4.5 C++.
`idlOrbix`	Generates IDL somewhat compatible with Orbix 2000 2.0 C++.
`idlDirectory <dir>`	Specifies the directory where IDL files will be created (default: target directory or JAR)
`idlMethodSignatures <>`	Specifies the method signatures used to trigger IDL code generation.
`iiop`	Generates CORBA stubs for EJBs.
`iiopDirectory <dir>`	Specifies the directory where IIOP stub files will be written (default: target directory or JAR)
`keepgenerated`	Keeps the generated `.java` files.
`librarydir`	Specifies a directory of shared Java EE libraries to add to the classpath. See "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.
`compiler <java.jdt>`	Selects the Java compiler to use. Defaults to JDT.
`debug`	Compiles debugging information into a class file.
`optimize`	Compiles with optimization on.
`nowarn`	Compiles without warnings.
`verbose`	Compiles with verbose output.
`deprecation`	Warns about deprecated calls.
`normi`	Passes flags through to Symantec's sj.
`runtimeflags`	Passes flags through to Java runtime
`classpath <path>`	Selects the classpath to use during compilation.
`clientJarOutputDir <dir>`	Specifies a directory to place generated client jar files. If not set, generated jar files are placed into the same directory location where the JVM is running.
`advanced`	Prints advanced usage options.

wlappc Ant Task Syntax

The basic syntax for using the `wlappc` Ant task determines the destination source directory location. This directory contains the files to be compiled by `wlappc`.

```
<wlappc source="${dest.dir}" />
```

The following is an example of a `wlappc` Ant task command that invokes two options (`idl` and `idlOrverWrite`) from Table 4-3.

```
<wlappc source="${dest.dir}"idl="true" idlOrverWrite="true" />
```

Syntax Differences between appc and wlappc

There are some syntax differences between `appc` and `wlappc`. For appc, the presence of a flag in the command is a boolean. For `wlappc`, the presence of a flag in the command means that the argument is required.

To illustrate, the following are examples of the same command, the first being an `appc` command and the second being a `wlappc` command:

```
java weblogic.appc -idl foo.ear
```

```
<wlappc source="${dest.dir} idl="true"/>
```

weblogic.appc Reference

The following sections describe how to use the command-line version of the `appc` compiler. The `weblogic.appc` command-line compiler reports any warnings or errors encountered in the descriptors and compiles all of the relevant modules into an EAR file, which can be deployed to WebLogic Server.

weblogic.appc Syntax

Use the following syntax to run `appc`:

```
prompt>java weblogic.appc [options] <ear, jar, or war file or directory>
```

weblogic.appc Options

The following are the available appc options:

Table 4-4 appc Options

Option	Description
-print	Prints the standard usage message.
-version	Prints appc version information.
-output <file>	Specifies an alternate output archive or directory. If not set, the output is placed in the source archive or directory.
-forceGeneration	Forces generation of EJB and JSP classes. Without this flag, the classes may not be regenerated (if determined to be unnecessary).
-library <file[[@name=<string>][@libspecver=<version>][@libimplver=<version|string>]]>	A comma-separated list of shared Java EE libraries. Optional name and version string information must be specified in the format described in "Referencing Shared Java EE Libraries in an Enterprise Application" on page 9-12.
-writeInferredDescriptors	Specifies that the application or module contains deployment descriptors with annotation information.
-lineNumbers	Adds line numbers to generated class files to aid in debugging.
-basicClientJar	Does not include deployment descriptors in client JARs generated for EJBs.
-idl	Generates IDL for EJB remote interfaces.
-idlOverwrite	Always overwrites existing IDL files.
-idlVerbose	Displays verbose information for IDL generation.
-idlNoValueTypes	Does not generate valuetypes and the methods/attributes that contain them.
-idlNoAbstractInterfaces	Does not generate abstract interfaces and methods/attributes that contain them.
-idlFactories	Generates factory methods for valuetypes.

Table 4-4 appc Options

`-idlVisibroker`	Generates IDL somewhat compatible with Visibroker 4.5 C++.
`-idlOrbix`	Generates IDL somewhat compatible with Orbix 2000 2.0 C++.
`-idlDirectory <dir>`	Specifies the directory where IDL files will be created (default: target directory or JAR)
`-idlMethodSignatures <>`	Specifies the method signatures used to trigger IDL code generation.
`-iiop`	Generates CORBA stubs for EJBs.
`-iiopDirectory <dir>`	Specifies the directory where IIOP stub files will be written (default: target directory or JAR)
`-keepgenerated`	Keeps the generated `.java` files.
`-compiler <javac>`	Selects the Java compiler to use.
`-g`	Compiles debugging information into a class file.
`-O`	Compiles with optimization on.
`-nowarn`	Compiles without warnings.
`-verbose`	Compiles with verbose output.
`-deprecation`	Warns about deprecated calls.
`-normi`	Passes flags through to Symantec's sj.
`-J<option>`	Passes flags through to Java runtime.
`-classpath <path>`	Selects the classpath to use during compilation.
`-clientJarOutputDir <dir>`	Specifies a directory to place generated client jar files. If not set, generated jar files are placed into the same directory location where the JVM is running.
`-advanced`	Prints advanced usage options.

Deploying and Packaging from a Split Development Directory

The following sections describe the steps for deploying WebLogic Server Java EE applications using the WebLogic split development directory environment:

Deploying Applications Using wldeploy

The `wldeploy` task provides an easy way to deploy directly from the split development directory. `wlcompile` provides most of the same arguments as the `weblogic.Deployer` directory. To deploy from a split development directory, you simply identify the build directory location as the deployable files, as in:

```
<wldeploy user="${user}" password="${password}"
       action="deploy" source="${dest.dir}"
       name="helloWorldEar" />
```

The above task is automatically created when you use `weblogic.BuildXMLGen` to create the `build.xml` file.

See for a complete command reference.

Packaging Applications Using wlpackage

The `wlpackage` Ant task uses the contents of both the source and build directories to create either a deployable archive file (`.EAR` file), or an exploded archive directory representing the Enterprise

Application (exploded `.EAR` directory). Use `wlpackage` when you want to deliver your application to another group or individual for evaluation, testing, performance profiling, or production deployment.

Archive versus Exploded Archive Directory

For production purposes, it is convenient to deploy Enterprise applications in exploded (unarchived) directory format. This applies also to stand-alone Web applications, EJBs, and connectors packaged as part of an Enterprise application. Using this format allows you to update files directly in the exploded directory rather than having to unarchive, edit, and rearchive the whole application. Using exploded archive directories also has other benefits, as described in Deployment Archive Files Versus Exploded Archive Directories in *Deploying Applications to WebLogic Server*.

You can also package applications in a single archived file, which is convenient for packaging modules and applications for distribution. Archive files are easier to copy, they use up fewer file handles than an exploded directory, and they can save disk space with file compression.

The Java classloader can search for Java class files (and other file types) in a JAR file the same way that it searches a directory in its classpath. Because the classloader can search a directory or a JAR file, you can deploy Java EE modules on WebLogic Server in either a JAR (archived) file or an exploded (unarchived) directory.

wlpackage Ant Task Example

In a production environment, use the `wlpackage` Ant task to package your split development directory application as a traditional EAR file that can be deployed to WebLogic Server. Continuing with the MedRec example, you would package your application as follows:

```
<wlpackage tofile="\physicianEAR\physicianEAR.ear"
           srcdir="\physicianEAR"
            destdir="\build\physicianEAR"/>

<wlpackage todir="\physicianEAR\explodedphysicianEar"
           srcdir="\src\physicianEAR"
            destdir="\build\physicianEAR" />
```

wlpackage Ant Task Attribute Reference

The following table describes the attributes of the wlpackage Ant task.

Table 5-1 Attributes of the wlpackage Ant Task

Attribute	Description	Data Type	Required?
tofile	Name of the EAR archive file into which the wlpackage Ant task packages the split development directory application.	String	You must specify one of the following two attributes: tofile or todir.
todir	Name of an exploded directory into which the wlpackage Ant task packages the split devleopment directory application.	String	You must specify one of the following two attributes: tofile or todir.
srcdir	Specifies the source directory of your split development directory application. The source directory contains all editable files for your project—Java source files, editable descriptor files, JSPs, static content, and so forth.	String	Yes.
destdir	Specifies the build directory of your split development directory application. It is assumed that you have already executed the wlcompile Ant task against the source directory to generate the needed components into the build directory; these components include compiled Java classes and generated deployment descriptors.	String	Yes.

Developing Applications for Production Redeployment

The following sections describes how to program and maintain applications use the production redeployment strategy:

What is Production Redeployment?

Production redeployment enables an Administrator to redeploy a new version of an application in a production environment without stopping the deployed application or otherwise interrupting the application's availability to clients. Production redeployment works by deploying a new version of an updated application alongside an older version of the same application. WebLogic Server automatically manages client connections so that only new client requests are directed to the new version. Clients already connected to the application during the redeployment continue to use the older, retiring version of the application until they complete their work.

See Using Production Redeployment to Upgrade Applications for more information.

Supported and Unsupported Application Types

Production redeployment only supports HTTP clients and RMI clients. Your development and design team must ensure that applications using production redeployment are not accessed by an unsupported client. WebLogic Server does not detect when unsupported clients access the application, and does not preserve unsupported client connections during production redeployment.

Enterprise Applications can contain any of the supported Java EE module types. Enterprise Applications can also include application-scoped JMS and JDBC modules.

If an Enterprise Application includes a JCA resource adapter module, the module:

- Must be JCA 1.5 compliant

- Must implement the `weblogic.connector.extensions.Suspendable` interface

- Must be used in an application-scoped manner, having `enable-access-outside-app` set to `false` (the default value).

Before resource adapters in a newer version of the EAR are deployed, resource adapters in the older application version receive a callback. WebLogic Server then deploys the newer application version and retires the entire older version of the EAR.

For a complete list of production redeployment requirements for resource adapters, see Production Redeployment in *Programming WebLogic Resource Adapters*.

Additional Application Support

Additional production redeployment support is provided for Enterprise Applications that are accessed by inbound JMS messages from a global JMS destination, and that use one or more message-driven beans as consumers. For this type of application, WebLogic Server suspends message-driven beans in the older, retiring application version before deploying message-driven beans in the newer version. Production redeployment is not supported with JMS consumers that use the JMS API for global JMS destinations. If the message-driven beans need to receive all messages published from topics, including messages published while bean are suspended, use durable subscribers.

Programming Requirements and Conventions

WebLogic Server performs production redeployment by deploying two instances of an application simultaneously. You must observe certain programming conventions to ensure that

multiple instances of the application can co-exist in a WebLogic Server domain. The following sections describe each programming convention required for using production redeployment.

Applications Should Be Self-Contained

As a best practice, applications that use the in-place redeployment strategy should be self-contained in their use of resources. This means you should generally use application-scoped JMS and JDBC resources, rather than global resources, whenever possible for versioned applications.

If an application must use a global resource, you must ensure that the application supports safe, concurrent access by multiple instances of the application. This same restriction also applies if the application uses external (separately-deployed) applications, or uses an external property file. WebLogic Server does not prevent the use of global resources with versioned applications, but you must ensure that resources are accessed in a safe manner.

Looking up a global JNDI resource from within a versioned application results in a warning message. To disable this check, set the JNDI environment property `weblogic.jndi.WLContext.ALLOW_GLOBAL_RESOURCE_LOOKUP` to `true` when performing the JNDI lookup.

Similarly, looking up an external application results in a warning unless you set the JNDI environment property, `weblogic.jndi.WLContext.ALLOW_EXTERNAL_APP_LOOKUP`, to `true`.

Versioned Applications Access the Current Version JNDI Tree by Default

WebLogic Server binds application-scoped resources, such as JMS and JDBC application modules, into a local JNDI tree available to the application. As with non-versioned applications, versioned applications can look up application-scoped resources directly from this local tree. Application-scoped JMS modules can be accessed via any supported JMS interfaces, such as the JMS API or a message-driven bean.

Application modules that are bound to the global JNDI tree should be accessed only from within the same application version. WebLogic Server performs version-aware JNDI lookups and bindings for global resources deployed in a versioned application. By default, an internal JNDI lookup of a global resource returns bindings for the same version of the application.

If the current version of the application cannot be found, you can use the JNDI environment property `weblogic.jndi.WLContext.RELAX_VERSION_LOOKUP` to return bindings from the currently active version of the application, rather than the same version.

WARNING: Set weblogic.jndi.WLContext.RELAX_VERSION_LOOKUP to true only if you are certain that the newer and older version of the resource that you are looking up are compatible with one another.

Security Providers Must Be Compatible

Any security provider used in the application must support the WebLogic Server application versioning SSPI. The default WebLogic Server security providers for authorization, role mapping, and credential mapping support the application versioning SSPI.

Applications Must Specify a Version Identifier

In order to use production redeployment, both the current, deployed version of the application and the updated version of the application must specify unique version identifiers. See "Assigning an Application Version" on page 6-5.

Applications Can Access Name and Identifier

Versioned applications can programmatically obtain both an application name, which remains constant across different versions, and an application identifier, which changes to provide a unique label for different versions of the application. Use the application name for basic display or error messages that refer to the application's name irrespective of the deployed version. Use the application ID when the application must provide unique identifier for the deployed version of the application. See "Accessing Version Information" on page 6-6 for more information about the MBean attributes that provide the name and identifier.

Client Applications Use Same Version when Possible

As described in "What is Production Redeployment?" on page 6-1, WebLogic Server attempts to route a client application's requests to the same version of the application until all of the client's in-progress work has completed. However, if an application version is retired using a timeout period, or is undeployed, the client's request will be routed to the active version of the application. In other words, a client's association with a given version of an application is maintained only on a "best-effort basis."

This behavior can be problematic for client applications that recursively access other applications when processing requests. WebLogic Server attempts to dispatch requests to the same versions of the recursively-accessed applications, but cannot guarantee that an intermediate application version is not undeployed manually or after a timeout period. If you have a group of related

applications with strict version requirements, Oracle recommends packaging all of the applications together to ensure version consistency during production redeployment.

Assigning an Application Version

Oracle recommends that you specify the version identifier in the MANIFEST.MF of the application, and automatically increment the version each time a new application is released for deployment. This ensures that production redeployment is always performed when the administrator or deployer redeploys the application.

For testing purposes, a deployer can also assign a version identifier to an application during deployment and redeployment. See Assigning a Version Identifier During Deployment and Redeployment in *Deploying Applications to WebLogic Server*.

Application Version Conventions

WebLogic Server obtains the application version from the value of the Weblogic-Application-Version property in the MANIFEST.MF file. The version string can be a maximum of 215 characters long, and must consist of valid characters as identified in Table 6-1.

Table 6-1 Valid and Invalid Characters

Valid ASCII Characters	Invalid Version Constructs
a-z	..
A-Z	.
0-9	
period ("."), underscore ("_"), or hyphen ("-") in combination with other characters	

For example, the following manifest file content describes an application with version "v920.beta":

```
Manifest-Version: 1.0
    Created-By: 1.4.1_05-b01 (Sun Microsystems Inc.)
    Weblogic-Application-Version: v920.beta
```

Upgrading Applications to Use Production Redeployment

If you are upgrading applications for deployment to WebLogic Server 9.2, note that the `Name` attribute retrieved from `AppDeploymentMBean` now returns a unique application identifier consisting of both the deployed application name and the application version string. Applications that require only the deployed application name must use the new `ApplicationName` attribute instead of the `Name` attribute. Applications that require a unique identifier can use either the `Name` or `ApplicationIdentifier` attribute, as described in "Accessing Version Information" on page 6-6.

Accessing Version Information

Your application code can use new MBean attributes to retrieve version information for display, logging, or other uses. The following table describes the read-only attributes provided by `ApplicationMBean`.

Table 6-2 Read-Only Version Attributes in ApplicationMBean

Attribute Name	Description
`ApplicationName`	A String that represents the deployment name of the application
`VersionIdentifier`	A String that uniquely identifies the current application version across all versions of the same application
`ApplicationIdentifier`	A String that uniquely identifies the current application version across all deployed applications and versions

`ApplicationRuntimeMBean` also provides version information in the new read-only attributes described in the following table.

Table 6-3 Read-Only Version Attributes in ApplicationRuntimeMBean

Attribute Name	Description
`ApplicationName`	A String that represents the deployment name of the application
`ApplicationVersion`	A string that represents the version of the application.
`ActiveVersionState`	An integer that indicates the current state of the active application version. Valid states for an active version are:
	• ACTIVATED—indicates that one or more modules of the application are active and available for processing new client requests.
	• PREPARED—indicates that WebLogic Server has prepared one or more modules of the application, but that it is not yet active.
	• UNPREPARED—indicates that no modules of the application are prepared or active.
	See the WebLogic Server 9.2 API Reference for more information.
	Note that the currently active version does not always correspond to the last-deployed version, because the Administrator can reverse the production redeployment process. See Rolling Back the Production Redeployment Process in *Deploying Applications to WebLogic Server*.

Using Java EE Annotations and Dependency Injection

The following sections describe the concepts of MetaData annotation and dependency injection:

- "Annotation Processing" on page 7-1

- "Dependency Injection of Resources" on page 7-3

- "Standard JDK Annotations" on page 7-4

- "Standard Security-Related JDK Annotations" on page 7-9

Annotation Processing

With Java EE annotations, the standard `application.xml` and `web.xml` deployment descriptors are optional. The Java EE programming model uses the JDK 5.0 annotations feature for Web containers, such as EJBs, servlets, Web applications, and JSPs.

Annotations simplify the application development process by allowing developers to specify within the Java class itself how the application component behaves in the container, requests for dependency injection, and so on. Annotations are an alternative to deployment descriptors that were required by older versions of Enterprise applications (J2EE 1.4 and earlier).

Annotation Parsing

The application components can use annotations to define their needs. Annotations reduce or eliminate the need to deal with deployment descriptors. Annotations simplify the development of application components. The deployment descriptor can still override values defined in the annotation. One usage of annotations is to define fields or methods that need Dependency

Injection (DI). Annotations are defined on the POJO (plain old Java object) component classes like the EJB or the servlet.

An annotation on a field or a method can declare that fields/methods need injection, as described in "Dependency Injection of Resources" on page 7-3. Annotations may also be applied to the class itself. The class-level annotations declare an entry in the application component's environment but do not cause the resource to be injected. Instead, the application component is expected to use JNDI or component context lookup method to lookup the entry. When the annotation is applied to the class, the JNDI name and the environment entry type must be specified explicitly.

Deployment View of Annotation Configuration

The Java EE Deployment API [JSR88] provides a way for developers to examine deployment descriptors. For example, consider an EJB Module that has no deployment descriptors. Assuming that it has some classes that have been declared as EJBs using annotations. A user of Session Helper will still be able to deal with the module as if it had the deployment descriptor. So the developer can modify the configuration information and it will be written out in a deployment plan. During deployment, such a plan will be honored and will override information from annotations.

Compiling Annotated Classes

The Weblogic Server utility `appc` (and its Ant equivalent `wlappc`) and `Appmerge` support metadata annotations. The `appmerge` and `appc` utilities take an application or module as inputs and process them to produce an output application or module respectively. When used with `-writeInferredDescriptors` flag, the output application/module will contain deployment descriptors with annotation information. The descriptors will also have the `metadata-complete` attribute set to `true`, as no annotation processing needs to be done if the output application or module is deployed directly. However, setting of `metadata-complete` attribute to `true` will also restrict `appmerge` and `appc` from processing annotations in case these tools are invoked on a previously processed application or module.

The original descriptors must be preserved in such cases to with an `.orig` suffix. If a developer wants to reapply annotation processing on the output application, they must restore the descriptors and use the `-writeInferredDescriptors` flag again. If `appmerge` or `appc` is used with `-writeInferredDescriptors` on an Enterprise application for which no standard deployment descriptor exists, the descriptor will be generated and written out based on the inference rules in the Java EE specification.

For more information on using `appc`, see "weblogic.appc Reference" on page 4-7. For more information on using `appmerge`, see "Using weblogic.appmerge to Merge Libraries" on page 9-19.

Dynamic Annotation Updates

Deployed modules can be updated using `update` deployment operation. If such an update has changes to deployment descriptor or updated classes, the container must consider annotation information again while processing the new deployment descriptor.

Containers use the descriptor framework's two-phase update mechanism to check the differences between the current and proposed descriptors. This mechanism also informs the containers about any changes in the non-dynamic properties. The containers then deal with such non-dynamic changes in their own specific ways. The container must perform annotation processing on the proposed descriptor to make sure that it is finding the differences against the right reference.

Similarly, some of the classes from a module could be updated during an update operation. If the container knows that these classes could affect configuration information through annotations, it makes sure that nothing has changed.

Dependency Injection of Resources

Dependency injection (DI) allows application components to declare dependencies on external resources and configuration parameters via annotations. The container reads these annotations and injects resources or environment entries into the application components. Dependency injection is simply an easier-to-program alternative to using the `javax` interfaces or JNDI APIs to look up resources.

A field or a method of an application component can be annotated with the `@Resource` annotation. Note that the container will unbox the environment entry as required to match it to a primitive type used for the injection field or method. Listing 7-1 illustrates how an application component uses the `@Resource` annotation to declare environment entries.

Listing 7-1

```
// fields

// The maximum number of tax exemptions, configured by the Deployer.
@Resource int maxExemptions;
```

```
        // The minimum number of tax exemptions, configured by the Deployer.
        @Resource int minExemptions;


        …..
        }
```

In the above code the `@Resource` annotation has not specified a name; therefore, the container would look for an `env-entry` name called `<class-name>/maxExemptions` and inject the value of that entry into the `maxExemptions` variable. The field or method may have any access qualifier (public, private, etc.). For all classes except application client main classes, the fields or methods must not be static. Because application clients use the same life cycle as J2EE applications, no instance of the application client main class is created by the application client container. Instead, the static main method is invoked. To support injection for the application client main class, the fields or methods annotated for injection must be static.

Application Life-Cycle Annotation Methods

An application component may need to perform initialization of its own after all resources have been injected. To support this case, one method of the class can be annotated with the `@PostConstruct` annotation. This method will be called after all injections have occurred and before the class is put into service. This method will be called even if the class doesn't request any resources to be injected. Similarly, for classes whose life cycle is managed by the container, the `@PreDestroy` annotation can be applied to one method that will be called when the class is taken out of service and will no longer be used by the container. Each class in a class hierarchy may have `@PostConstruct` and `@PreDestroy` methods.

The order in which the methods are called matches the order of the class hierarchy, with methods on a superclass being called before methods on a subclass. From the Java EE side only the application client container is involved in invoking these life-cycle methods for Java EE clients. The life cycle methods for Java EE clients must be static. The Java EE client just supports the `@PostConstruct` callback.

Standard JDK Annotations

This section provides reference information about the following annotations:

-

- "javax.annotation.PreDestroy" on page 7-6
- "javax.annotation.Resource" on page 7-7
- "javax.annotation.Resources" on page 7-8

For detailed information about EJB-specific annotations for WebLogic Server Enterprise JavaBeans, see *Programming WebLogic Enterprise JavaBeans, Version 3.0*.

For detailed information about web component-specific annotations WebLogic Server applications, see "WebLogic Annotation for Web Components" in *Developing Web Applications, Servlets, and JSPs for WebLogic Server.*

javax.annotation.PostConstruct

Description

Target: Method

Specifies the life-cycle callback method that the application component should execute before the first business method invocation and after dependency injection is done to perform any initialization. This method will be called after all injections have occurred and before the class is put into service. This method will be called even if the class doesn't request any resources to be injected.

You must specify a `@PostConstruct` method in any component that includes dependency injection.

Only one method in the component can be annotated with this annotation.

The method annotated with `@PostConstruct` must follow these requirements:

- The method must not have any parameters, except in the case of EJB interceptors, in which case it takes an `javax.interceptor.InvocationContext` object as defined by the EJB specification.
- The return type of the method must be `void`.
- The method must not throw a checked exception.
- The method may be `public`, `protected`, `package private` or `private`.
- The method must not be `static` except for the application client.
- The method may be `final` or `non-final`, except in the case of EJBs where it must be `non-final`.

- If the method throws an unchecked exception the class must not be put into service. In the case of EJBs, the method annotated with `PostConstruct` can handle exceptions and cleanup before the bean instance is discarded.

This annotation does not have any attributes.

javax.annotation.PreDestroy

Description

Target: Method

Specifies the life-cycle callback method that signals that the application component is about to be destroyed by the container. You typically apply this annotation to methods that release resources that the class has been holding.

Only one method in the bean class can be annotated with this annotation.

The method annotated with `@PreDestroy` must follow these requirements:

- The method must not have any parameters, except in the case of EJB interceptors, in which case it takes an `javax.interceptor.InvocationContext` object as defined by the EJB specification.

- The return type of the method must be `void`.

- The method must not throw a checked exception.

- The method may be `public`, `protected`, `package private` or `private`.

- The method must not be `static` except for the application client.

- The method may be `final` or `non-final`, except in the case of EJBs where it must be `non-final`.

- If the method throws an unchecked exception the class must not be put into service. In the case of EJBs, the method annotated with `PreDestroy` can handle exceptions and cleanup before the bean instance is discarded.

This annotation does not have any attributes.

javax.annotation.Resource

Description

Target: Class, Method, Field

Specifies a dependence on an external resource, such as a JDBC data source or a JMS destination or connection factory.

If you specify the annotation on a field or method, the application component injects an instance of the requested resource into the bean when the bean is initialized. If you apply the annotation to a class, the annotation declares a resource that the component will look up at runtime.

Attributes

Table 0-1 Attributes of the javax.annotation.Resource Annotation

Name	Description	Data Type	Required?
name	Specifies the JNDI name of the resource.	String	No
	If you apply the `@Resource` annotation to a field, the default value of the `name` attribute is the field name, qualified by the class name. If you apply it to a method, the default value is the component property name corresponding to the method, qualified by the class name. If you apply the annotation to class, there is no default value and thus you are required to specify the attribute.		
type	Specifies the Java data type of the resource.	Class	No
	If you apply the `@Resource` annotation to a field, the default value of the `type` attribute is the type of the field. If you apply it to a method, the default is the type of the component property. If you apply it to a class, there is no default value and thus you are required to specify this attribute.		
authentication Type	Specifies the authentication type to use for the resource.	Authentica tionType	No
	Valid values for this attribute are:		
	• `AuthenticationType.CONTAINER`		
	• `AuthenticationType.APPLICATION`		
	Default value is `AuthenticationType.CONTAINER`		

Table 0-1 Attributes of the javax.annotation.Resource Annotation

Name	Description	Data Type	Required?
shareable	Indicates whether a resource can be shared between this component and other components. Valid values for this attribute are `true` and `false`. Default value is `true`.	boolean	No

mappedName	Specifies a WebLogic Server-specific name to which the component reference should be mapped.
	
	However, if you do not specify a JNDI name in the WebLogic deployment descriptor file, then the value of `mappedName` will always be used as the JNDI name to look up. For example:
	
	`@Resource(mappedName = "http://www.bea.com";)`
	`URL url;`
	
	`@Resource(mappedName="customerDB")`
	`DataSource db;`
	
	`@Resource(mappedName =`
	`"jms/ConnectionFactory")`
	`ConnectionFactory connectionFactory;`
	
	`@Resource(mappedName = "jms/Queue")`
	`Queue queue;`
	
	In other words, `MappedName` is honored as JNDI name only when there is no JNDI name specified elsewhere, typically in the WebLogic deployment descriptor file.
		String	No

description	Specifies a description of the resource.	String	No

javax.annotation.Resources

Description

Target: Class

Specifies an array of `@Resource` annotations. Since repeated annotations are not allowed, the Resources annotation acts as a container for multiple resource declarations.

Attributes

Table 0-2 Attributes of the javax.annotation.Resources Annotation

Name	Description	Data Type	Required?
value	Specifies the array of `@Resource` annotations.	`Resourc e[]`	Yes

Standard Security-Related JDK Annotations

This section provides reference information about the following annotations:

- "javax.annotation.security.DeclareRoles" on page 7-9

- "javax.annotation.security.DenyAll" on page 7-10

- "javax.annotation.security.PermitAll" on page 7-10

- "javax.annotation.security.RolesAllowed" on page 7-10

- "javax.annotation.security.RunAs" on page 7-11

javax.annotation.security.DeclareRoles

Description

Target: Class

Defines the security roles that will be used in the Java EE container.

You typically use this annotation to define roles that can be tested from within the methods of the annotated class, such as using the `isUserInRole` method. You can also use the annotation to explicitly declare roles that are implicitly declared if you use the `@RolesAllowed` annotation on the class or a method of the class.

You create security roles in WebLogic Server using the Administration Console. For details, see Manage Security Roles.

Attributes

Table 0-3 Attributes of the javax.annotation.security.DeclareRoles Annotation

Name	Description	Data Type	Required?
value	Specifies an array of security roles that will be used in the Java EE container.	String[]	Yes

javax.annotation.security.DenyAll

Description

Target: Method

Specifies that no security role is allowed to access the annotated method, or in other words, the method is excluded from execution in the Java EE container.

This annotation does not have any attributes.

javax.annotation.security.PermitAll

Description

Target: Method

Specifies that all security roles currently defined for WebLogic Server are allowed to access the annotated method.

This annotation does not have any attributes.

javax.annotation.security.RolesAllowed

Description

Target: Class, Method

Specifies the list of security roles that are allowed to access methods in the Java EE container.

If you specify it at the class-level, then it applies to all methods in the application component. If you specify it at the method-level, then it only applies to that method. If you specify the annotation at both the class- and method-level, the method value overrides the class value.

You create security roles in WebLogic Server using the Administration Console. For details, see Manage Security Roles.

Attributes

Table 0-4 Attributes of the javax.annotation.security.RolesAllowed Annotation

Name	Description	Data Type	Required?
value	List of security roles that are allowed to access methods of the Java EE container.	String[]	Yes

javax.annotation.security.RunAs

Description

Target: Class

Specifies the security role which actually executes the Java EE container.

The security role must exist in the WebLogic Server security realm and map to a user or group. For details, see Manage Security Roles.

Attributes

Table 0-5 Attributes of the javax.annotation.security.RunAs Annotation

Name	Description	Data Type	Required?
value	Specifies the security role that the Java EE container should run as.	String	Yes

Understanding WebLogic Server Application Classloading

The following sections provide an overview of Java classloaders, followed by details about WebLogic Server Java EE application classloading.

Java Classloading

Classloaders are a fundamental module of the Java language. A classloader is a part of the Java virtual machine (JVM) that loads classes into memory; a classloader is responsible for finding and loading class files at run time. Every successful Java programmer needs to understand classloaders and their behavior. This section provides an overview of Java classloaders.

Java Classloader Hierarchy

Classloaders contain a hierarchy with parent classloaders and child classloaders. The relationship between parent and child classloaders is analogous to the object relationship of super classes and subclasses. The bootstrap classloader is the root of the Java classloader hierarchy. The Java virtual machine (JVM) creates the bootstrap classloader, which loads the Java development kit

(JDK) internal classes and `java.*` packages included in the JVM. (For example, the bootstrap classloader loads `java.lang.String`.)

The extensions classloader is a child of the bootstrap classloader. The extensions classloader loads any JAR files placed in the extensions directory of the JDK. This is a convenient means to extending the JDK without adding entries to the classpath. However, anything in the extensions directory must be self-contained and can only refer to classes in the extensions directory or JDK classes.

The system classpath classloader extends the JDK extensions classloader. The system classpath classloader loads the classes from the classpath of the JVM. Application-specific classloaders (including WebLogic Server classloaders) are children of the system classpath classloader.

Note: What Oracle refers to as a "system classpath classloader" is often referred to as the "application classloader" in contexts outside of WebLogic Server. When discussing classloaders in WebLogic Server, Oracle uses the term "system" to differentiate from classloaders related to Java EE applications or libraries (which Oracle refers to as "application classloaders").

Loading a Class

Classloaders use a delegation model when loading a class. The classloader implementation first checks its cache to see if the requested class has already been loaded. This class verification improves performance in that its cached memory copy is used instead of repeated loading of a class from disk. If the class is not found in its cache, the current classloader asks its parent for the class. Only if the parent cannot load the class does the classloader attempt to load the class. If a class exists in both the parent and child classloaders, the parent version is loaded. This delegation model is followed to avoid multiple copies of the same form being loaded. Multiple copies of the same class can lead to a `ClassCastException`.

Classloaders ask their parent classloader to load a class before attempting to load the class themselves. Classloaders in WebLogic Server that are associated with Web applications can be configured to check locally first before asking their parent for the class. This allows Web applications to use their own versions of third-party classes, which might also be used as part of the WebLogic Server product. The "prefer-web-inf-classes Element" on page 8-2 section discusses this in more detail.

prefer-web-inf-classes Element

The `weblogic.xml` Web application deployment descriptor contains a `<prefer-web-inf-classes>` element (a sub-element of the `<container-descriptor>`

element). By default, this element is set to False. Setting this element to True subverts the classloader delegation model so that class definitions from the Web application are loaded in preference to class definitions in higher-level classloaders. This allows a Web application to use its own version of a third-party class, which might also be part of WebLogic Server. See "weblogic.xml Deployment Descriptor Elements."

When using this feature, you must be careful not to mix instances created from the Web application's class definition with issuances created from the server's definition. If such instances are mixed, a ClassCastException results.

Listing 8-1 illustrates the prefer-web-inf-classes element, its description and default value.

Listing 8-1   prefer-web-inf-classes Element

```
/**

* If true, classes located in the WEB-INF directory of a web-app will be

* loaded in preference to classes loaded in the application or system

* classloader.

* @default false

*/

boolean isPreferWebInfClasses();

void setPreferWebInfClasses(boolean b);
```

Changing Classes in a Running Program

WebLogic Server allows you to deploy newer versions of application modules such as EJBs while the server is running. This process is known as hot-deploy or hot-redeploy and is closely related to classloading.

Java classloaders do not have any standard mechanism to undeploy or unload a set of classes, nor can they load new versions of classes. In order to make updates to classes in a running virtual machine, the classloader that loaded the changed classes must be replaced with a new classloader. When a classloader is replaced, all classes that were loaded from that classloader (or any classloaders that are offspring of that classloader) must be reloaded. Any instances of these classes must be re-instantiated.

In WebLogic Server, each application has a hierarchy of classloaders that are offspring of the system classloader. These hierarchies allow applications or parts of applications to be individually reloaded without affecting the rest of the system. "WebLogic Server Application Classloading" on page 8-4 discusses this topic.

WebLogic Server Application Classloading

The following sections provide an overview of the WebLogic Server application classloaders:

- "Overview of WebLogic Server Application Classloading" on page 8-4

- "Application Classloader Hierarchy" on page 8-5

- "Custom Module Classloader Hierarchies" on page 8-7

- "Individual EJB Classloader for Implementation Classes" on page 8-12

- "Application Classloading and Pass-by-Value or Reference" on page 8-14

- "Using a Filtering Classloader" on page 8-14

Overview of WebLogic Server Application Classloading

WebLogic Server classloading is centered on the concept of an application. An application is normally packaged in an Enterprise Archive (EAR) file containing application classes. Everything within an EAR file is considered part of the same application. The following may be part of an EAR or can be loaded as stand-alone applications:

- An Enterprise JavaBean (EJB) JAR file

- A Web application WAR file

- A resource adapter RAR file

Notes:

- For information on Resource Adapters and classloading, see "About Resource Adapter Classes" on page 8-18.

- For information on overriding generic application files while classloading, see Generic File Loading Overrides in *Deploying Applications to WebLogic Server*.

If you deploy an EJB and a Web application separately, they are considered two applications. If they are deployed together within an EAR file, they are one application. You deploy modules together in an EAR file for them to be considered part of the same application.

Every application receives its own classloader hierarchy; the parent of this hierarchy is the system classpath classloader. This isolates applications so that application A cannot see the classloaders or classes of application B. In hierarchy classloaders, no sibling or friend concepts exist. Application code only has visibility to classes loaded by the classloader associated with the application (or module) and classes that are loaded by classloaders that are ancestors of the application (or module) classloader. This allows WebLogic Server to host multiple isolated applications within the same JVM.

Application Classloader Hierarchy

WebLogic Server automatically creates a hierarchy of classloaders when an application is deployed. The root classloader in this hierarchy loads any EJB JAR files in the application. A child classloader is created for each Web application WAR file.

Because it is common for Web applications to call EJBs, the WebLogic Server application classloader architecture allows JavaServer Page (JSP) files and servlets to see the EJB interfaces in their parent classloader. This architecture also allows Web applications to be redeployed without redeploying the EJB tier. In practice, it is more common to change JSP files and servlets than to change the EJB tier.

The following graphic illustrates this WebLogic Server application classloading concept.

Figure 8-1   WebLogic Server Classloading
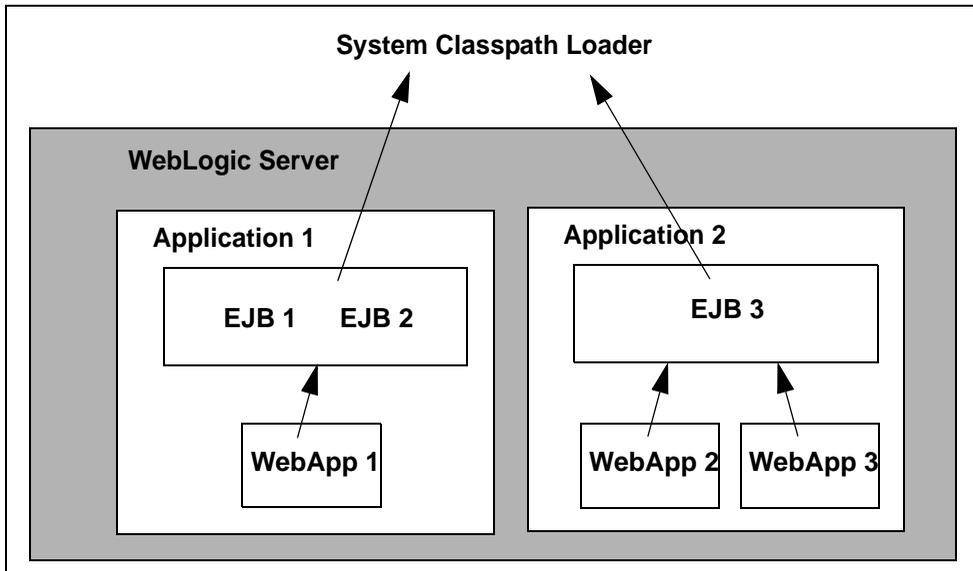
If your application includes servlets and JSPs that use EJBs:

- Package the servlets and JSPs in a WAR file

- Package the Enterprise JavaBeans in an EJB JAR file

- Package the WAR and JAR files in an EAR file

- Deploy the EAR file

Although you could deploy the WAR and JAR files separately, deploying them together in an
EAR file produces a classloader arrangement that allows the servlets and JSPs to find the EJB
classes. If you deploy the WAR and JAR files separately, WebLogic Server creates sibling
classloaders for them. This means that you must include the EJB home and remote interfaces in
the WAR file, and WebLogic Server must use the RMI stub and skeleton classes for EJB calls,
just as it does when EJB clients and implementation classes are in different JVMs. This concept
is discussed in more detail in the next section "Application Classloading and Pass-by-Value or
Reference" on page 8-14.

Note: The Web application classloader contains all classes for the Web application except for the JSP class. The JSP class obtains its own classloader, which is a child of the Web application classloader. This allows JSPs to be individually reloaded.

Custom Module Classloader Hierarchies

You can create custom classloader hierarchies for an application allowing for better control over class visibility and reloadability. You achieve this by defining a `classloader-structure` element in the `weblogic-application.xml` deployment descriptor file.

The following diagram illustrates how classloaders are organized by default for WebLogic applications. An application level classloader exists where all EJB classes are loaded. For each Web module, there is a separate child classloader for the classes of that module.

For simplicity, JSP classloaders are not described in the following diagram.

Figure 8-2 Standard Classloader Hierarchy
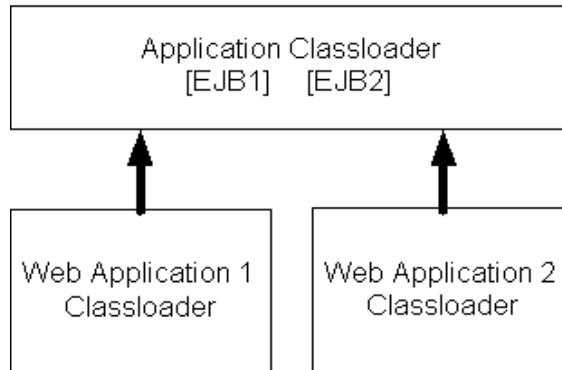
This hierarchy is optimal for most applications, because it allows call-by-reference semantics when you invoke EJBs. It also allows Web modules to be independently reloaded without affecting other modules. Further, it allows code running in one of the Web modules to load classes from any of the EJB modules. This is convenient, as it can prevent a Web module from including the interfaces for EJBs that is uses. Note that some of those benefits are not strictly Java EE-compliant.

The ability to create custom module classloaders provides a mechanism to declare alternate classloader organizations that allow the following:

- Reloading individual EJB modules independently

- Reloading groups of modules to be reloaded together

- Reversing the parent child relationship between specific Web modules and EJB modules

- Namespace separation between EJB modules

Declaring the Classloader Hierarchy

You can declare the classloader hierarchy in the WebLogic-specific application deployment descriptor `weblogic-application.xml`.

The DTD for this declaration is as follows:

Listing 8-2   Declaring the Classloader Hierarchy

```
<!ELEMENT classloader-structure (module-ref*, classloader-structure*)>

<!ELEMENT module-ref (module-uri)>

<!ELEMENT module-uri (#PCDATA)>
```

The top-level element in `weblogic-application.xml` includes an optional `classloader-structure` element. If you do not specify this element, then the standard classloader is used. Also, if you do not include a particular module in the definition, it is assigned a classloader, as in the standard hierarchy. That is, EJB modules are associated with the application Root classloader, and Web application modules have their own classloaders.

The `classloader-structure` element allows for the nesting of `classloader-structure` stanzas, so that you can describe an arbitrary hierarchy of classloaders. There is currently a limitation of three levels. The outermost entry indicates the application classloader. For any modules not listed, the standard hierarchy is assumed.

Note: JSP classloaders are not included in this definition scheme. JSPs are always loaded into a classloader that is a child of the classloader associated with the Web module to which it belongs.

For more information on the DTD elements, refer to Appendix A, "Enterprise Application Deployment Descriptor Elements."

The following is an example of a classloader declaration (defined in the
`classloader-structure` element in `weblogic-application.xml`):

Listing 8-3   Example Classloader Declaration

```
<classloader-structure>

      <module-ref>

            <module-uri>ejb1.jar</module-uri>

      </module-ref>

      <module-ref>

            <module-uri>web3.war</module-uri>

      </module-ref>


      <classloader-structure>

            <module-ref>

                  <module-uri>web1.war</module-uri>

            </module-ref>

      </classloader-structure>


      <classloader-structure>

            <module-ref>

                  <module-uri>ejb3.jar</module-uri>

            </module-ref>

            <module-ref>

                  <module-uri>web2.war</module-uri>

            </module-ref>

            <classloader-structure>

                  <module-ref>
```

```
                                <module-uri>web4.war</module-uri>

                        </module-ref>

                </classloader-structure>

                <classloader-structure>

                        <module-ref>

                                <module-uri>ejb2.jar</module-uri>

                        </module-ref>

                </classloader-structure>

        </classloader-structure>

</classloader-structure>
```

The organization of the nesting indicates the classloader hierarchy. The above stanza leads to a hierarchy shown in the following diagram.

Figure 8-3   Example Classloader Hierarchy
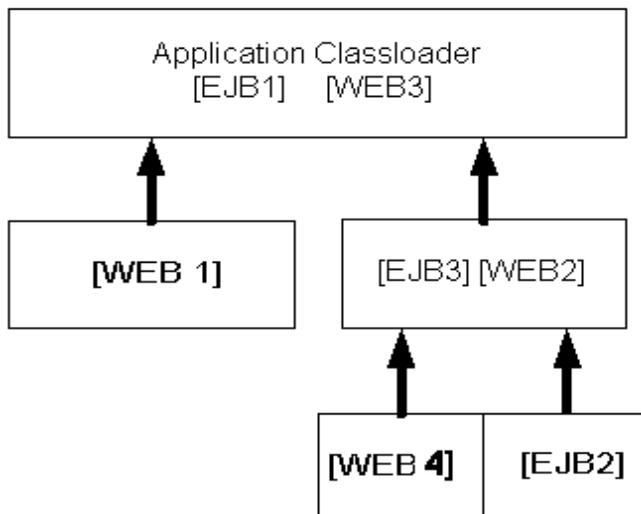
User-Defined Classloader Restrictions

User-defined classloader restrictions give you better control over what is reloadable and provide inter-module class visibility. This feature is primarily for developers. It is useful for iterative development, but the reloading aspect of this feature is not recommended for production use, because it is possible to corrupt a running application if an update includes invalid elements. Custom classloader arrangements for namespace separation and class visibility are acceptable for production use. However, programmers should be aware that the Java EE specifications say that applications should not depend on any given classloader organization.

Some classloader hierarchies can cause modules within an application to behave more like modules in two separate applications. For example, if you place an EJB in its own classloader so that it can be reloaded individually, you receive call-by-value semantics rather than the call-by-reference optimization Oracle provides in our standard classloader hierarchy. Also note that if you use a custom hierarchy, you might end up with stale references. Therefore, if you reload an EJB module, you should also reload calling modules.

There are some restrictions to creating user-defined module classloader hierarchies; these are discussed in the following sections.

Servlet Reloading Disabled

If you use a custom classloader hierarchy, servlet reloading is disabled for Web applications in that particular application.

Nesting Depth

Nesting is limited to three levels (including the application classloader). Deeper nestings lead to a deployment exception.

Module Types

Custom classloader hierarchies are currently restricted to Web and EJB modules.

Duplicate Entries

Duplicate entries lead to a deployment exception.

Interfaces

The standard WebLogic Server classloader hierarchy makes EJB interfaces available to all modules in the application. Thus other modules can invoke an EJB, even though they do not include the interface classes in their own module. This is possible because EJBs are always

loaded into the root classloader and all other modules either share that classloader or have a classloader that is a child of that classloader.

With the custom classloader feature, you can configure a classloader hierarchy so that a callee's classes are not visible to the caller. In this case, the calling module must include the interface classes. This is the same requirement that exists when invoking on modules in a separate application.

Call-by-Value Semantics

The standard classloader hierarchy provided with WebLogic Server allows for calls between modules within an application to use call-by-reference semantics. This is because the caller is always using the same classloader or a child classloader of the callee. With this feature, it is possible to configure the classloader hierarchy so that two modules are in separate branches of the classloader tree. In this case, call-by-value semantics are used.

In-Flight Work

Be aware that the classloader switch required for reloading is not atomic across modules. In fact, updates to applications in general are not atomic. For this reason, it is possible that different in-flight operations (operations that are occurring while a change is being made) might end up accessing different versions of classes depending on timing.

Development Use Only

The development-use-only feature is intended for development use. Because updates are not atomic, this feature is not suitable for production use.

Individual EJB Classloader for Implementation Classes

WebLogic Server allows you to reload individual EJB modules without requiring you to reload other modules at the same time and having to redeploy the entire EJB module. This feature is similar to how JSPs are currently reloaded in the WebLogic Server servlet container.

Because EJB classes are invoked through an interface, it is possible to load individual EJB implementation classes in their own classloader. This way, these classes can be reloaded individually without having to redeploy the entire EJB module. Below is a diagram of what the classloader hierarchy for a single EJB module would look like. The module contains two EJBs (Foo and Bar). This would be a sub-tree of the general application hierarchy described in the previous section.

Figure 8-4 Example Classloader Hierarchy for a Single EJB Module
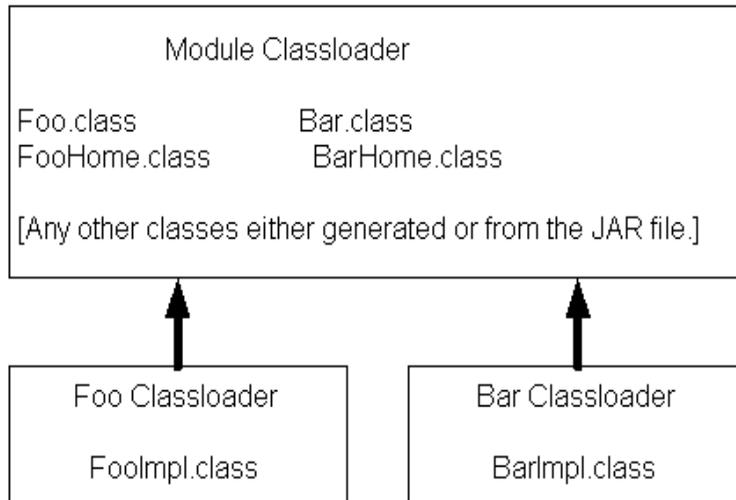
To perform a partial update of files relative to the root of the exploded application, use the following command line:

Listing 8-4 Performing a Partial File Update

```
java weblogic.Deployer -adminurl url -user user -password password

-name myapp -redeploy myejb/foo.class
```

After the `-redeploy` command, you provide a list of files relative to the root of the exploded application that you want to update. This might be the path to a specific element (as above) or a module (or any set of elements and modules). For example:

Listing 8-5 Providing a List of Relative Files for Update

```
java weblogic.Deployer -adminurl url -user user -password password

-name myapp -redeploy mywar myejb/foo.class anotherejb
```

Given a set of files to be updated, the system tries to figure out the minimum set of things it needs to redeploy. Redeploying only an EJB `impl` class causes only that class to be redeployed. If you specify the whole EJB (in the above example, `anotherejb`) or if you change and update the EJB home interface, the entire EJB module must be redeployed.

Depending on the classloader hierarchy, this redeployment may lead to other modules being redeployed. Specifically, if other modules share the EJB classloader or are loaded into a classloader that is a child to the EJB's classloader (as in the WebLogic Server standard classloader module) then those modules are also reloaded.

Application Classloading and Pass-by-Value or Reference

Modern programming languages use two common parameter passing models: pass-by-value and pass-by-reference. With pass-by-value, parameters and return values are copied for each method call. With pass-by-reference, a pointer (or reference) to the actual object is passed to the method. Pass by reference improves performance because it avoids copying objects, but it also allows a method to modify the state of a passed parameter.

WebLogic Server includes an optimization to improve the performance of Remote Method Interface (RMI) calls within the server. Rather than using pass by value and the RMI subsystem's marshalling and unmarshalling facilities, the server makes a direct Java method call using pass by reference. This mechanism greatly improves performance and is also used for EJB 2.0 local interfaces.

RMI call optimization and call by reference can only be used when the caller and callee are within the same application. As usual, this is related to classloaders. Because applications have their own classloader hierarchy, any application class has a definition in both classloaders and receives a ClassCastException error if you try to assign between applications. To work around this, WebLogic Server uses call-by-value between applications, even if they are within the same JVM.

Note: Calls between applications are slower than calls within the same application. Deploy modules together as an EAR file to enable fast RMI calls and use of the EJB 2.0 local interfaces.

Using a Filtering Classloader

In WebLogic Server, any JAR file present in the system classpath is loaded by the WebLogic Server system classloader. All applications running within a server instance are loaded in application classloaders which are children of the system classloader. In this implementation of the system classloader, applications cannot use different versions of third-party jars which are

already present in the system classloader. Every child classloader asks the parent (the system classloader) for a particular class and cannot load classes which are seen by the parent.

For example, if a class called `com.foo.Baz` exists in both `$CLASSPATH` as well as the application EAR, then the class from the `$CLASSPATH` is loaded and not the one from the EAR. Since `weblogic.jar` is in the `$CLASSPATH`, applications can not override any WebLogic Server classes.

The following sections define and describe how to use a filtering classloader:

- "What is a Filtering ClassLoader" on page 8-15
- "Configuring a FilteringClassLoader" on page 8-15
- "Resource Loading Order" on page 8-16

What is a Filtering ClassLoader

The `FilteringClassLoader` provides a mechanism for you to configure deployment descriptors to explicitly specify that certain packages should always be loaded from the application, rather than being loaded by the system classloader. This allows you to use alternate versions of applications such as Xerces and Ant.

The `FilteringClassLoader` sits between the application classloader and the system. It is a child of the system classloader and the parent of the application classloader. The `FilteringClassLoader` intercepts the `loadClass(String className)` method and compares the `className` with a list of packages specified in `weblogic-application.xml` file. If the package matches the `className`, the `FilteringClassLoader` throws a `ClassNotFoundException`. This exception notifies the application classloader to load this class from the application.

Configuring a FilteringClassLoader

To configure the `FilteringClassLoader` to specify a certain package is loaded from an application, add a `prefer-application-packages` descriptor element to the `weblogic-application.xml` which details the list of packages to be loaded from the application. The following example specifies that `org.apache.log4j.*` and `antlr.*` packages are loaded from the application, not the system classloader:

```
<prefer-application-packages>

  <package-name>org.apache.log4j.*</package-name>

  <package-name>antlr.*</package-name>
```

```
</prefer-application-packages>
```

Resource Loading Order

The resource loading order is the order in which `java.lang.ClassLoader` methods `getResource()` and `getResources()` return resources. When filtering is enabled, this order is slightly different from the case when filtering is disabled. Filtering is enabled implies that there are one or more package patterns in the `FilteringClassLoader`. Without any filtering (default), the resources are collected in the top-down order of the classloader tree. For instance, if Web (1) requests resources, the resources are grouped in the following order: Sys (3), App (2) and Web(1). See Figure 8-5.

Figure 8-5   Using the System Classloader

```
System (3)
  |
App (2)
  |
Web (1)
```
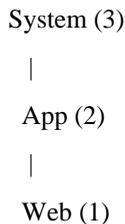
To be more explicit, given a resource `/META-INF/foo.xml` which exists in all the classloaders, would return the following list of URLs:

```
META-INF/foo.xml - from the System ClassLoader (3)

META-INF/foo.xml - from the App ClassLoader (2)

META-INF/foo.xml - from the Web ClassLoader (1)
```

When filtering is enabled, the resources from the child of the `FilteringClassLoader` (an application classloader) down to the calling classloader are returned before the ones from the system classloader. In Figure 8-6, if the same resource existed in all the classloaders (D), (B) and (A) one would get them in the following order if requested by the Web classloader:

```
META-INF/foo.xml - from the App ClassLoader (B)

META-INF/foo.xml - from the Web ClassLoader (A)

META-INF/foo.xml - from the System ClassLoader (D)
```

Note: The resources are returned in the default Java EE delegation model beneath the `FilteringClassLoader`. Only the resources from the parent of the `FilteringClassLoader` are appended to the end of the enumeration being returned.

Figure 8-6   Using a Filtering Classloading Implementation

System (D)

 |

FilteringClassLoader (filterList := x.y.*) (C)

 |

App (B)

 |

Web (A)
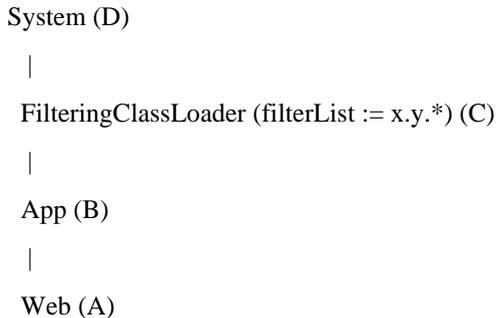
If the application classloader requested the same resource, the following order would be obtained.

```
META-INF/foo.xml - from the App ClassLoader (B)

META-INF/foo.xml - from the System ClassLoader (D)
```

For `getResource()`, only the first descriptor is returned and `getResourceAsStream()` returns the `inputStream` of the first resource.

Resolving Class References Between Modules and Applications

Your applications may use many different Java classes, including Enterprise Beans, servlets and JavaServer Pages, utility classes, and third-party packages. WebLogic Server deploys applications in separate classloaders to maintain independence and to facilitate dynamic redeployment and undeployment. Because of this, you need to package your application classes in such a way that each module has access to the classes it depends on. In some cases, you may have to include a set of classes in more than one application or module. This section describes how WebLogic Server uses multiple classloaders so that you can stage your applications successfully.

About Resource Adapter Classes

Each resource adapter now uses its own classloader to load classes (similar to Web applications). As a result, modules like Web applications and EJBs that are packaged along with a resource adapter in an application archive (EAR file) do not have visibility into the resource adapter's classes. If such visibility is required, you must place the resource adapter classes in `APP-INF/classes`. You can also archive these classes (using the JAR utility) and place them in the `APP-INF/lib` of the application archive.

Make sure that no resource-adapter specific classes exist in your WebLogic Server system classpath. If you need to use resource adapter-specific classes with Web modules (for example, an EJB or Web application), you must bundle these classes in the corresponding module's archive file (for example, the JAR file for EJBs or the WAR file for Web applications).

Packaging Shared Utility Classes

WebLogic Server provides a location within an EAR file where you can store shared utility classes. Place utility JAR files in the `APP-INF/lib` directory and individual classes in the `APP-INF/classes` directory. (Do not place JAR files in the `/classes` directory or classes in the `/lib` directory.) These classes are loaded into the root classloader for the application.

This feature obviates the need to place utility classes in the system classpath or place classes in an EJB JAR file (which depends on the standard WebLogic Server classloader hierarchy). Be aware that using this feature is subtly different from using the manifest `Class-Path` described in the following section. With this feature, class definitions are shared across the application. With manifest `Class-Path`, the classpath of the referencing module is simply extended, which means that separate copies of the classes exist for each module.

Manifest Class-Path

The Java EE specification provides the manifest `Class-Path` entry as a means for a module to specify that it requires an auxiliary JAR of classes. You only need to use this manifest `Class-Path` entry if you have additional supporting JAR files as part of your EJB JAR or WAR file. In such cases, when you create the JAR or WAR file, you must include a manifest file with a `Class-Path` element that references the required JAR files.

The following is a simple manifest file that references a `utility.jar` file:

```
Manifest-Version: 1.0 [CRLF]
Class-Path: utility.jar [CRLF]
```

In the first line of the manifest file, you must always include the `Manifest-Version` attribute, followed by a new line (CR | LF |CRLF) and then the `Class-Path` attribute. More information about the manifest format can be found at:

http://java.sun.com/j2se/1.4/docs/guide/jar/jar.html#JAR

The manifest `Class-Path` entries refer to other archives relative to the current archive in which these entries are defined. This structure allows multiple WAR files and EJB JAR files to share a common library JAR. For example, if a WAR file contains a manifest entry of `y.jar`, this entry should be next to the WAR file (not within it) as follows:

/<directory>/x.war

/<directory>/y.jars

The manifest file itself should be located in the archive at `META-INF/MANIFEST.MF`.

For more information, see

http://java.sun.com/docs/books/tutorial/jar/basics/manifest.html.

Sharing Applications and Modules By Using Java EE Libraries

Java EE libraries provide an easy way to share one or more different types of Java EE modules among multiple Enterprise Applications. A Java EE library is a single module or collection of modules that is registered with the Java EE application container upon deployment. For more information, see Chapter 9, "Creating Shared Java EE Libraries and Optional Packages."

Adding JARs to the System Classpath

WebLogic Server includes a `lib` subdirectory, located in the domain directory, that you can use to add one or more JAR files to the WebLogic Server system classpath when servers start up. The `lib` subdirectory is intended for JAR files that change infrequently and are required by all or most applications deployed in the server, or by WebLogic Server itself. For example, you might use the `lib` directory to store third-party utility classes that are required by all deployments in a domain. You can also use it to apply patches to WebLogic Server.

The `lib` directory is not recommended as a general-purpose method for sharing a JARs between one or two applications deployed in a domain, or for sharing JARs that need to be updated periodically. If you update a JAR in the `lib` directory, you must reboot all servers in the domain in order for applications to realize the change. If you need to share a JAR file or Java EE modules

among several applications, use the Java EE libraries feature described in "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.

To share JARs using the `lib` directory:

1. Shutdown all servers in the domain.

2. Copy the JAR file(s) to share into a `lib` subdirectory of the domain directory. For example:

```
mkdir c:\bea\wlserver_10.3\samples\domains\wl_server\lib
cp c:\3rdpartyjars\utility.jar
   c:\bea\wlserver_10.3\samples\domains\wl_server\lib
```

   Note: WebLogic Server must have read access to the `lib` directory during startup.

   Note: The Administration Server does not automatically copy files in the `lib` directory to Managed Servers on remote machines. If you have Managed Servers that do not share the same physical domain directory as the Administration Server, you must manually copy JAR file(s) to the *domain_name*/`lib` directory on the Managed Server machines.

3. Start the Administration Server and all Managed Servers in the domain. WebLogic Server appends JAR files found in the `lib` directory to the system classpath. Multiple files are added in alphabetical order.

Creating Shared Java EE Libraries and Optional Packages

The following sections describe how to share components and classes among applications using shared Java EE libraries and optional packages:

Overview of Shared Java EE Libraries and Optional Packages

The shared Java EE library feature in WebLogic Server provides an easy way to share one or more different types of Java EE modules among multiple Enterprise Applications. A shared Java EE library is a single module or collection of modules that is registered with the Java EE application container upon deployment. A shared Java EE library can be any of the following:

- standalone EJB module

- standalone Web application module

- multiple EJB modules packaged in an Enterprise Application

- multiple Web application modules package in an Enterprise Application

- single plain JAR file

Oracle recommends that you package a shared Java EE library into its appropriate archive file (EAR, JAR, or WAR). However, for development purposes, you may choose to deploy shared Java EE libraries as exploded archive directories to facilitate repeated updates and redeployments.

After the shared Java EE library has been registered, you can deploy Enterprise Applications that reference the library. Each referencing application receives a reference to the required library on deployment, and can use the modules that make up the library as if they were packaged as part of the referencing application itself. The library classes are added to the classpath of the referencing application, and the referencing application's deployment descriptors are merged (in memory) with those of the modules that make up the shared Java EE library.

In general, this topic discusses shared Java EE libraries that can be referenced only by Enterprise Applications. You can also create libraries that can be referenced only by another Web application. The functionality is very similar to application libraries, although the method of referencing them is slightly different. See "Web Application Shared Java EE Library Information" on page 9-22 for details.

Note: WebLogic Server also provides a simple way to add one or more JAR files to the WebLogic Server System classpath, using the `lib` subdirectory of the domain directory. See "Adding JARs to the System Classpath" on page 8-19.

Optional Packages

WebLogic Server supports optional packages as described in the Java EE 5.0 Specification, Section 8.2 Optional Package Support, with versioning described in Optional Package Versioning. Optional packages provide similar functionality to Java EE libraries, allowing you to easily share a single JAR file among multiple applications. As with Java EE libraries, optional packages must first be registered with WebLogic Server by deploying the associated JAR file as an optional package. After registering the package, you can deploy Java EE modules that reference the package in their manifest files.

Optional packages are also supported as Java EE shared libraries in `weblogic.BuildXMLGen`, whereby all manifests of an application and its modules are scanned to look for optional package references. If optional package references are found they are added to the `wlcompile` and `appc` tasks in the generated `build.xml` file.

Optional packages differ from Java EE libraries because optional packages can be referenced from any Java EE module (EAR, JAR, WAR, or RAR archive) or exploded archive directory. Java EE libraries can be referenced only from a valid Enterprise Application.

For example, third-party Web Application Framework classes needed by multiple Web Applications can be packaged and deployed in a single JAR file, and referenced by multiple Web Application modules in the domain. Optional packages, rather than Java EE libraries, are used in this case, because the individual Web Application modules must reference the shared JAR file. (With Java EE libraries, only a complete Enterprise Application can reference the library).

Note: Oracle documentation and WebLogic Server utilities use the term *library* to refer to both Java EE libraries and optional packages. Optional packages are called out only when necessary.

Library Directories

The Java EE platform provides several mechanisms for applications to use optional packages and shared libraries. Libraries can be bundled with an application or may be installed separately for use by any application. An EAR file may contain a directory that contains libraries packaged in JAR files. The `library-directory` element of the EAR file's deployment descriptor contains the name of this directory. If a `library-directory` element isn't specified, or if the EAR file does not contain a deployment descriptor, the directory named `lib` is used. An empty `library-directory` element may be used to specify that there is no library directory. All files in this directory (but not in subdirectories) with a `.jar` extension must be made available to all components packaged in the EAR file, including application clients. These libraries may reference other libraries, either bundled with the application or installed separately.

This feature is similar to the `APP-INF/lib` feature supported in WebLogic Server. If both `APP-INF/lib` and `library-directory` exist, then the jars in the `library-directory` would take precedence; that is, they would be placed before the `APP-INF/lib` jar files in the classpath. For more information on `APP-INF/lib`, see "Resolving Class References Between Modules and Applications" on page 8-17 and "Organizing Shared Classes in a Split Development Directory" on page 3-11.

Versioning Support for Libraries

WebLogic Server supports versioning of shared Java EE libraries, so that referencing applications can specify a required minimum version of the library to use, or an exact, required version. WebLogic Server supports two levels of versioning for shared Java EE libraries, as described in the Optional Package Versioning document:

- Specification Version—Identifies the version number of the specification (for example, the Java EE specification version) to which a shared Java EE library or optional package conforms.

- Implementation Version—Identifies the version number of the actual code implementation for the library or package. For example, this would correspond to the actual revision number or release number of your code. Note that you must also provide a specification version in order to specify an implementation version.

As a best practice, Oracle recommends that you always include version information (an implementation version, or both an implementation and specification version) when creating shared Java EE libraries. Creating and updating version information as you develop shared components allows you to deploy multiple versions of those components simultaneously for testing. If you include no version information, or fail to increment the version string, then you must undeploy existing libraries before you can deploy the newer one. See "Deploying Shared Java EE Libraries and Dependent Applications" on page 9-21.

Versioning information in the referencing application determines the library and package version requirements for that application. Different applications can require different versions of a given library or package. For example, a production application may require a specific version of a library, because only that library has been fully approved for production use. An internal application may be configured to always use a minimum version of the same library. Applications that require no specific version can be configured to use the latest version of the library. "Referencing Shared Java EE Libraries in an Enterprise Application" on page 9-12.

Shared Java EE Libraries and Optional Packages Compared

Optional packages and shared Java EE libraries have the following features in common:

- Both are registered with WebLogic Server instances at deployment time.

- Both support an optional implementation version and specification version string.

- Applications that reference shared Java EE libraries and optional packages can specify required versions for the shared files.

- Optional packages can reference other optional packages, and shared Java EE libraries can reference other shared Java EE libraries.

Optional packages differ from shared Java EE Libraries in the following basic ways:

- Optional packages are plain JAR files, whereas shared Java EE libraries can be plain JAR files, Java EE Enterprise Applications, or standalone Java EE modules (EJB and Web applications). This means that libraries can have valid Java EE and WebLogic Server deployment descriptors. Any deployment descriptors in an optional package JAR file are ignored.

- Any Java EE application or module can reference an optional package (using `META-INF/MANIFEST.MF`), whereas only Enterprise Applications and Web applications can reference a shared Java EE library (using `weblogic-application.xml` or `weblogic.xml`)

In general, use shared Java EE libraries when you need to share one or more EJB, Web Application or Enterprise Application modules among different Enterprise Applications. Use optional packages when you need to share one or more classes (packaged in a JAR file) among different Java EE modules.

Plain JAR files can be shared either as libraries or optional packages. Use optional packages if you want to:

- Share a plain JAR file among multiple Java EE modules

- Reference shared JAR files from other shared JARs

- Share plain JARs as described by the Java EE 5.0 specification

Use shared Java EE libraries to share a plain JAR file if you only need to reference the JAR file from one or more Enterprise Applications, and you do not need to maintain strict compliance with the Java EE specification.

> **Note:** Oracle documentation and WebLogic Server utilities use the term *shared Java EE library* to refer to both libraries and optional packages. Optional packages are called out only when necessary.

Additional Information

For information about deploying and managing shared Java EE libraries, optional packages, and referencing applications from the Administrator's perspective, see Deploying Shared Java EE Libraries and Dependent Applications in *Deploying Applications to WebLogic Server*.

Creating Shared Java EE Libraries

To create a new shared Java EE library that you can share with multiple applications:

1. Assemble the shared Java EE library into a valid, deployable Java EE module or Enterprise Application. The library must have the required Java EE deployment descriptors for the Java EE module or for an Enterprise Application.

    See "Assembling Shared Java EE Library Files" on page 9-6.

2. Assemble optional package classes into a working directory.

    See "Assembling Optional Package Class Files" on page 9-7.

3. Create and edit the MANIFEST.MF file for the shared Java EE library to specify the name and version string information.

    See "Editing Manifest Attributes for Shared Java EE Libraries" on page 9-9.

4. Package the shared Java EE library for distribution and deployment.

    See "Packaging Shared Java EE Libraries for Distribution and Deployment" on page 9-11.

Assembling Shared Java EE Library Files

The following types of Java EE modules can be deployed as a shared Java EE library:

- An EJB module, either an exploded directory or packaged in a JAR file.

- A Web Application module, either an exploded directory or packaged in a WAR file.

- An Enterprise application, either an exploded directory or packaged in an EAR file.

- A plain Java class or classes packaged in a JAR file.

- A shared Java EE library referenced from another library. (See "Web Application Shared Java EE Library Information" on page 9-22.)

Shared Java EE libraries have the following restrictions:

- You must ensure that context roots in Web application modules of the shared Java EE library do not conflict with context roots in the referencing Enterprise Application. If necessary, you can configure referencing applications to override a library's context root. See "Referencing Shared Java EE Libraries in an Enterprise Application" on page 9-12.

- Shared Java EE libraries cannot be nested. For example, if you are deploying an EAR as a shared Java EE library, the entire EAR must be designated as the library. You cannot designate individual Java EE modules within the EAR as separate, named libraries.

- As with any other Java EE module or Enterprise Application, a shared Java EE library must be configured for deployment to the target servers or clusters in your domain. This means that a library requires valid Java EE deployment descriptors as well as WebLogic Server-specific deployment descriptors and an optional deployment plan. See *Deploying Applications to WebLogic Server*.

Oracle recommends packaging shared Java EE libraries as Enterprise Applications, rather than as standalone Java EE modules. This is because the URI of a standalone module is derived from the deployment name, which can change depending on how the module is deployed. By default, WebLogic Server uses the deployment archive filename or exploded archive directory name as the deployment name. If you redeploy a standalone shared Java EE library from a different file or location, the deployment name and URI also change, and referencing applications that use the wrong URI cannot access the deployed library.

If you choose to deploy a shared Java EE library as a standalone Java EE module, always specify a known deployment name during deployment and use that name as the URI in referencing applications.

Assembling Optional Package Class Files

Any set of classes can be organized into an optional package file. The collection of shared classes will eventually be packaged into a standard JAR archive. However, because you will need to edit the manifest file for the JAR, begin by assembling all class files into a working directory:

1. Create a working directory for the new optional package. For example:

   ```
   mkdir /apps/myOptPkg
   ```

2. Copy the compiled class files into the working directory, creating the appropriate package sudirectories as necessary. For example:

```
mkdir -p /apps/myOptPkg/org/myorg/myProduct
cp /build/classes/myOptPkg/org/myOrg/myProduct/*.class
/apps/myOptPkg/org/myOrg/myProduct
```

3. If you already have a JAR file that you want to use as an optional package, extract its contents into the working directory so that you can edit the manifest file:

```
cd /apps/myOptPkg
jar xvf /build/libraries/myLib.jar
```

Editing Manifest Attributes for Shared Java EE Libraries

The name and version information for a shared Java EE library are specified in the `META-INF/MANIFEST.MF` file. Table 9-1 describes the valid shared Java EE library manifest attributes.

Table 9-1 Manifest Attributes for Java EE Libraries

Attribute	Description
Extension-Name	An optional string value that identifies the name of the shared Java EE library. Referencing applications must use the exact `Extension-Name` value to use the library.
	As a best practice, always specify an `Extension-Name` value for each library. If you do not specify an extension name, one is derived from the deployment name of the library. Default deployment names are different for archive and exploded archive deployments, and they can be set to arbitrary values in the deployment command.

Table 9-1 Manifest Attributes for Java EE Libraries

Attribute	Description
Specification-Version	An optional String value that defines the specification version of the shared Java EE library. Referencing applications can optionally specify a required Specification-Version for a library; if the exact specification version is not available, deployment of the referencing application fails.
	The Specification-Version uses the following format:
	Major/minor version format, with version and revision numbers separated by periods (such as "9.0.1.1")
	Referencing applications can be configured to require either an exact version of the shared Java EE library, a minimum version, or the latest available version.
	The specification version for a shared Java EE library can also be set at the command-line when deploying the library, with some restrictions. See "Deploying Shared Java EE Libraries and Dependent Applications" on page 9-21.
Implementation-Version	An optional String value that defines the code implementation version of the shared Java EE library. You can provide an Implementation-Version only if you have also defined a Specification-Version.
	Implementation-Version uses the following formats:
	• Major/minor version format, with version and revision numbers separated by periods (such as "9.0.1.1")
	• Text format, with named versions (such as "9011Beta" or "9.0.1.1.B")
	If you use the major/minor version format, referencing applications can be configured to require either an exact version of the shared Java EE library, a minimum version, or the latest available version. If you use the text format, referencing applications must specify the exact version of the library.
	The implementation version for a shared Java EE library can also be set at the command-line when deploying the library, with some restrictions. See "Deploying Shared Java EE Libraries and Dependent Applications" on page 9-21.

To specify attributes in a manifest file:

1. Open (or create) the manifest file using a text editor. For the example shared Java EE library, you would use the commands:

```
cd /apps/myLibrary
mkdir META-INF
emacs META-INF/MANIFEST.MF
```

For the optional package example, use:

```
cd /apps/myOptPkg
mkdir META-INF
emacs META-INF/MANIFEST.MF
```

2. In the text editor, add a string value to specify the name of the shared Java EE library. For example:

```
Extension-Name: myExtension
```

Applications that reference the library must specify the exact `Extension-Name` in order to use the shared files.

3. As a best practice, enter the optional version information for the shared Java EE library. For example:

```
Extension-Name: myExtension
Specification-Version: 2.0
Implementation-Version: 9.0.0
```

Using the major/minor format for the version identifiers provides the most flexibility when referencing the library from another application (see Table)

Note: Although you can optionally specify the `Specification-Version` and `Implementation-Version` at the command-line during deployment, Oracle recommends that you include these strings in the `MANIFEST.MF` file. Including version strings in the manifest ensures that you can deploy new versions of the library alongside older versions. See "Deploying Shared Java EE Libraries and Dependent Applications" on page 9-21.

Packaging Shared Java EE Libraries for Distribution and Deployment

If you are delivering the shared Java EE Library or optional package for deployment by an Administrator, package the deployment files into an archive file (an `.EAR` file or standalone module archive file for shared Java EE libraries, or a simple `.JAR` file for optional packages) for distribution. See "Deploying and Packaging from a Split Development Directory" on page 5-1.

Because a shared Java EE library is packaged as a standard Java EE application or standalone module, you may also choose to export a library's deployment configuration to a deployment plan, as described in *Deploying Applications to WebLogic Server*. Optional package `.JAR` files contain no deployment descriptors and cannot be exported.

For development purposes, you may choose to deploy libraries as exploded archive directories to facilitate repeated updates and redeployments.

Referencing Shared Java EE Libraries in an Enterprise Application

A Java EE application can reference a registered shared Java EE library using entries in the application's `weblogic-application.xml` deployment descriptor. Table 9-2 describes the XML elements that define a library reference.

Table 9-2 weblogic-application.xml Elements for Referencing a Shared Java EE Library

Element	Description
library-ref	`library-ref` is the parent element in which you define a reference to a shared Java EE library. Enclose all other elements within `library-ref`.
library-name	A required string value that specifies the name of the shared Java EE library to use. `library-name` must exactly match the value of the `Extension-Name` attribute in the library's manifest file. (See Table 9-4.)
specification-version	An optional String value that defines the required specification version of the shared Java EE library. If this element is not set, the application uses a matching library with the highest specification version. If you specify a string value using major/minor version format, the application uses a matching library with the highest specification version that is not below the configured value. If all available libraries are below the configured `specification-version`, the application cannot be deployed. The required version can be further constrained by using the `exact-match` element, described below.
	If you specify a String value that does not use major/minor versioning conventions (for example, 9.2BETA) the application requires a shared Java EE library having the exact same string value in the `Specification-Version` attribute in the library's manifest file. (See Table 9-4.)

Table 9-2 weblogic-application.xml Elements for Referencing a Shared Java EE Library

Element	Description
implementation-version	An optional String value that specifies the required implementation version of the shared Java EE library. If this element is not set, the application uses a matching library with the highest implementation version. If you specify a string value using major/minor version format, the application uses a matching library with the highest implementation version that is not below the configured value. If all available libraries are below the configured implementation-version, the application cannot be deployed. The required implementation version can be further constrained by using the exact-match element, described below.
	If you specify a String value that does not use major/minor versioning conventions (for example, 9.2BETA) the application requires a shared Java EE library having the exact same string value in the Implementation-Version attribute in the library's manifest file. (See Table 9-4.)
exact-match	An optional boolean value that determines whether the application should use a shared Java EE library with a higher specification or implementation version than the configured value, if one is available. By default this element is false, which means that WebLogic Server uses higher-versioned libraries if they are available. Set this element to true to require the exact matching version as specified in the specification-version and implementation-version elements.
context-root	An optional String value that provides an alternate context root to use for a Web application shared Java EE library. Use this element if the context root of a library conflicts with the context root of a Web Application in the referencing Java EE application.
	Web application shared Java EE library refers to special kind of library: a Web application that is referenced by another Web application. See "Web Application Shared Java EE Library Information" on page 9-22.

For example, this simple entry in the weblogic-application.xml descriptor references a shared Java EE library, myLibrary:

```
<library-ref>
   <library-name>myLibrary</library-name>
</library-ref>
```

In the above example, WebLogic Server attempts to find a library name `myLibrary` when deploying the dependent application. If more than one copy of `myLibrary` is registered, WebLogic Server selects the library with the highest specification version. If multiple copies of the library use the selected specification version, WebLogic Server selects the copy having the highest implementation version.

This example references a shared Java EE library with a requirement for the specification version:

```
<library-ref>
    <library-name>myLibrary</library-name>
    <specification-version>2.0</specification-version>
</library-ref>
```

In the above example, WebLogic Server looks for matching libraries having a specification version of 2.0 or higher. If multiple libraries are at or above version 2.0, WebLogic Server examines the selected libraries that use Float values for their implementation version and selects the one with the highest version. Note that WebLogic Server ignores any selected libraries that have a non-Float value for the implementation version.

This example references a shared Java EE library with both a specification version and a non-Float value implementation version:

```
<library-ref>
    <library-name>myLibrary</library-name>
    <specification-version>2.0</specification-version>
    <implementation-version>81Beta</implementation-version>
</library-ref>
```

In the above example, WebLogic Server searches for a library having a specification version of 2.0 or higher, and having an exact match of `81Beta` for the implementation version.

The following example requires an exact match for both the specification and implementation versions:

```
<library-ref>
    <library-name>myLibrary</library-name>
    <specification-version>2.0</specification-version>
    <implementation-version>8.1</implementation-version>
    <exact-match>true</exact-match>
</library-ref>
```

The following example specifies a `context-root-override`, which in turn, refers to the old `context-root` specified in one of its libraries and the new `context-root` that should be used instead. (override):

```
<library-ref>
    <library-name>myLibrary</library-name>
    <library-context-root-override>
    <library-context-root>webapp</library-context-root>
    <override-value>mywebapp</override-value>
    </library-context-root-override>
</library-ref>
```

Overriding context-roots Within a Referenced Enterprise Library

A Java EE application can override `context-roots` within a referenced EAR library using entries in the application's `weblogic-application.xml` deployment descriptor. Table 9-3 describes the XML elements that override `context-roots` in a library reference.

Table 9-3 weblogic-application.xml Elements for Overriding a Shared Java EE Library

Element	Description
`library-context-root`	An optional String value that overrides the `context-root` elements declared in libraries. In the absence of this element, the library's `context-root` is used.
	Only a referencing application (i.e., a user application) can override the `context-root` elements declared in its libraries.
`override-value`	An optional String value that specifies the value of the `library-context-root-override` element when overriding the context-root elements declared in libraries. In the absence of these elements, the library's `context-root` is used.

The following example specifies a `context-root-override`, which in turn, refers to the old `context-root` specified in one of its libraries and the new `context-root` that should be used instead. (override):

```
<library-ref>
    <library-name>myLibrary</library-name>
```

```
    <specification-version>2.0</specification-version>
    <implementation-version>8.1</implementation-version>
    <exact-match>true</exact-match>
</library-ref>

<library-context-root-override>
    <library-context-root>webapp</library-context-root>
    <override-value>mywebapp</override-value>
</library-context-root-override>
```

In the above example, the current application refers to `myLibrary`, which contains a Web application with a `context-root` of `webapp`. The only way to override this reference is to declare a `library-context-root-override` that maps `webapp` to `mywebapp`.

URIs for Shared Java EE Libraries Deployed As a Standalone Module

When referencing the URI of a shared Java EE library that was deployed as a standalone module (EJB or Web Application), note that the module URI corresponds to the deployment name of the shared Java EE library. This can be a name that was manually assigned during deployment, the name of the archive file that was deployed, or the name of the exploded archive directory that was deployed. If you redeploy the same module using a different file name or from a different location, the default deployment name also changes and referencing applications must be updated to use the correct URI.

To avoid this problem, deploy all shared Java EE libraries as Enterprise Applications, rather than as standalone modules. If you choose to deploy a library as a standalone Java EE module, always specify a known deployment name and use that name as the URI in referencing applications.

Referencing Optional Packages from a Java EE Application or Module

Any Java EE archive (JAR, WAR, RAR, EAR) can reference one or more registered optional packages using attributes in the archive's manifest file.

Table 9-4 Manifest Attributes for Referencing Optional Packages

Attribute	Description
Extension-List *logical_name* [...]	A required String value that defines a logical name for an optional package dependency. You can use multiple values in the Extension-List attribute to designate multiple optional package dependencies. For example: `Extension-List: dependency1 dependency2`
[*logical_name*-]Extension-Name	A required string value that identifies the name of an optional package dependency. This value must match the Extension-Name attribute defined in the optional package's manifest file. If you are referencing multiple optional packages from a single archive, prepend the appropriate logical name to the Extension-Name attribute. For example: `dependency1-Extension-Name: myOptPkg`

Table 9-4 Manifest Attributes for Referencing Optional Packages

Attribute	Description
[*logical_name*-]Specification-Version	An optional String value that defines the required specification version of an optional package. If this element is not set, the archive uses a matching package with the highest specification version. If you include a specification-version value using the major/minor version format, the archive uses a matching package with the highest specification version that is not below the configured value. If all available package are below the configured specification-version, the archive cannot be deployed.
	If you specify a String value that does not use major/minor versioning conventions (for example, 9.2BETA) the archive requires a matching optional package having the exact same string value in the Specification-Version attribute in the package's manifest file. (See Table on page 9.)
	If you are referencing multiple optional packages from a single archive, prepend the appropriate logical name to the Specification-Version attribute.
[*logical_name*-]Implementation-Version	An optional String value that specifies the required implementation version of an optional package. If this element is not set, the archive uses a matching package with the highest implementation version. If you specify a string value using the major/minor version format, the archive uses a matching package with the highest implementation version that is not below the configured value. If all available libraries are below the configured implementation-version, the application cannot be deployed.
	If you specify a String value that does not use major/minor versioning conventions (for example, 9.2BETA) the archive requires a matching optional package having the exact same string value in the Implementation-Version attribute in the package's manifest file. (See Table on page 9.)
	If you are referencing multiple optional packages from a single archive, prepend the appropriate logical name to the Implementation-Version attribute.

For example, this simple entry in the manifest file for a dependent archive references two optional packages, myAppPkg and my3rdPartyPkg:

```
Extension-List: internal 3rdparty
internal-Extension-Name: myAppPkg
3rdparty-Extension-Name: my3rdPartyPkg
```

This example requires a specification version of 2.0 or higher for `myAppPkg`:

```
Extension-List: internal 3rdparty
internal-Extension-Name: myAppPkg
3rdparty-Extension-Name: my3rdPartyPkg
internal-Specification-Version: 2.0
```

This example requires a specification version of 2.0 or higher for `myAppPkg`, and an exact match for the implementation version of `my3rdPartyPkg`:

```
Extension-List: internal 3rdparty
internal-Extension-Name: myAppPkg
3rdparty-Extension-Name: my3rdPartyPkg
internal-Specification-Version: 2.0
3rdparty-Implementation-Version: 8.1GA
```

By default, when WebLogic Server deploys an application or module and it cannot resolve a reference in the application's manifest file to an optional package, WebLogic Server prints a warning, but continues with the deployment anyway. You can change this behavior by setting the system property `weblogic.application.RequireOptionalPackages` to `true` when you start WebLogic Server, either at the command line or in the command script file from which you start the server. Setting this system property to `true` means that WebLogic Server does *not* attempt to deploy an application or module if it cannot resolve an optional package reference in its manifest file.

Using weblogic.appmerge to Merge Libraries

`weblogic.appmerge` is a tool that is used to merge libraries into an application, with merged contents and merged descriptors. It also has the ability to write a merged application to disk. You can then use `weblogic.appmerge` to understand a library merge by examining the merged application you have written to disk.

- "Using weblogic.appmerge from the CLI" on page 9-20
- "Using weblogic.appmerge as an Ant Task" on page 9-20

Using weblogic.appmerge from the CLI

Invoke weblogic.appmerge using the following syntax:

```
java weblogic.appmerge [options] <ear, jar, war file, or directory>
```

where valid options are shown in Table 9-5:

Table 9-5 weblogic.appmerge Options

Option	Comment
-help	Print the standard usage message.
-version	Print version information.
-output <file>	Specifies an alternate output archive or directory. If not set, output is placed in the source archive or directory.
-plan <file>	Specifies an optional deployment plan.
-verbose	Provide more verbose output.

-library <file>	Comma-separated list of libraries. Each library may optionally set its name and versions, if not already set in its manifest, using the following syntax:
	
	`<file> [@name=<string>@libspecver=<version>`
	`@libimplver=<version|string>].`

-librarydir <dir>	Registers all files in specified directory as libraries.
-writeInferredDescriptors	Specifies that the application or module contains deployment descriptors with annotation information.

Example:

```
$ java weblogic.appmerge -output CompleteSportsApp.ear -library Weather
.war,Calendar.ear SportsApp.ear
```

Using weblogic.appmerge as an Ant Task

The ant task provides similar functionality as the command line utility. It supports `source`, `output`, `libraryDir`, `plan` and `verbose` attributes as well as multiple `<library>` sub-elements. Here is an example:

```
<taskdef name="appmerge" classname="weblogic.ant.taskdefs.j2ee.AppMergeTask"/>
<appmerge source="SportsApp.ear" output="CompleteSportsApp.ear">
   <library file="Weather.war"/>
   <library file="Calendar.ear"/>
</appmerge>
```

Integrating Shared Java EE Libraries with the Split Development Directory Environment

The `BuildXMLGen` includes a `-librarydir` option to generate build targets that include one or more shared Java EE library directories. See "Generating a Basic build.xml File Using weblogic.BuildXMLGen" on page 3-13.

The `wlcompile` and `wlappc` Ant tasks include a `librarydir` attribute and `library` element to specify one or more shared Java EE library directories to include in the classpath for application builds. See "Building Applications in a Split Development Directory" on page 4-1.

Deploying Shared Java EE Libraries and Dependent Applications

Shared Java EE libraries are registered with one or more WebLogic Server instances by deploying them to the target servers and indicating that the deployments are to be shared. Shared Java EE libraries must be targeted to the same WebLogic Server instances you want to deploy applications that reference the libraries. If you try to deploy a referencing application to a server instance that has not registered a required library, deployment of the referencing application fails. See Registering Libraries with WebLogic Server in *Deploying Applications to WebLogic Server* for more information.

See Install a Java EE Library for detailed instructions on installing (deploying) a shared Java EE library using the Administration Console. See Target a Shared Java EE Library to a Server or Cluster for instructions on using the Administration Console to target the library to the server or cluster to which the application that is referencing the library is also targeted.

If you use the `wldeploy` Ant task as part of your iterative development process, use the `library`, `libImplVer`, and `libSpecVer` attributes to deploy a shared Java EE library. See Appendix B, "wldeploy Ant Task Reference," for details and examples.

After registering a shared Java EE library, you can deploy applications and archives that depend on the library. Dependent applications can be deployed only if the target servers have registered all required libraries, and the registered deployments meet the version requirements of the application or archive. See Deploying Applications that Reference Libraries in *Deploying Applications to WebLogic Server* for more information.

Web Application Shared Java EE Library Information

In general, this topic discusses shared Java EE libraries that can be referenced only by Enterprise Applications. You can also create libraries that can be referenced only by another Web application. The functionality is very similar to application libraries, although the method of referencing them is slightly different.

Note: For simplicity, this section uses the term *Web application library* when referring to a shared Java EE library that is referenced only by another Web application.

In particular:

- Web application libraries can only be referenced by other Web applications.

- Rather than update the `weblogic-application.xml` file, Web applications reference Web application libraries by updating the `weblogic.xml` deployment descriptor file. The elements are almost same as those described in "Referencing Shared Java EE Libraries in an Enterprise Application" on page 9-12; the only difference is that the `<context-root>` child element of `<library-ref>` is ignored in this case.

- You cannot reference any other type of shared Java EE library (EJB, Enterprise application, or plain JAR file) from the `weblogic.xml` deployment descriptor file of a Web Application.

Other than these differences in how they are referenced, the way to create, package, and deploy a Web application library is the same as that of a standard shared Java EE library.

Using WebApp Libraries With Web Applications

Just as standard shared Java EE applications can be deployed to WebLogic Server as `application-libraries`, a standard Web application can be deployed to WebLogic Server as a `webapp-library` so that other Web applications can refer to these libraries.

Web application libraries facilitate the reuse of code and resources. Such libraries also help you separate out third-party Web applications or frameworks that your Web application might be using. Furthermore, common resources can be packaged separately as libraries and referenced in

different Web applications, so that you don't have to bundle them with each Web application. When you include a `webapp-library` in your Web application, at deployment time the container merges all the static resources, classes, and `JAR` files into your Web application.

The first step in using a WebApp library is to register a Web application as a `webapp-library`. This can be accomplished by deploying a Web application using either the Administration Console or the `weblogic.Deployer` tool as a library. To make other Web applications refer to this library, their `weblogic.xml` file must have a `library-ref` element pointing to the `webapp-library`, as follows:

```
<library-ref>
<library-name>BaseWebApp</library-name>
<specification-version>2.0</specification-version>
<implementation-version>8.1beta</implementation-version>
<exact-match>false</exact-match>
</library-ref>
```

When multiple libraries are present, the `CLASSPATH/resource` path precedence order follows the order in which the `library-refs` elements appear in the `weblogic.xml` file.

Accessing Registered Shared Java EE Library Information with LibraryRuntimeMBean

Each deployed shared Java EE library is represented by a `LibraryRuntimeMBean`. You can use this MBean to obtain information about the library itself, such as its name or version. You can also obtain the `ApplicationRuntimeMBeans` associated with deployed applications. `ApplicationRuntimeMBean` provides two methods to access the libraries that the application is using:

- `getLibraryRuntimes()` returns the shared Java EE libraries referenced in the `weblogic-application.xml` file.

- `getOptionalPackageRuntimes()` returns the optional packages referenced in the manifest file.

For more information, see the WebLogic Server API Reference.

Order of Precedence of Modules When Referencing Shared Java EE Libraries

When an Enterprise Application references one or more shared Java EE libraries, and the application is deployed to WebLogic Server, the server internally merges the information in the `weblogic-application.xml` file of the referencing Enterprise Application with the information in the deployment descriptors of the referenced libraries. The order in which this happens is as follows:

1. When the Enterprise Application is deployed, WebLogic Server reads its `weblogic-application.xml` deployment descriptor.

2. WebLogic Server reads the deployment descriptors of any referenced shared Java EE libraries. Depending on the type of library (Enterprise Application, EJB, or Web application), the read file might be `weblogic-application.xml`, `weblogic.xml`, `weblogic-ejb-jar.xml`, and so on.

3. WebLogic Server first merges the referenced shared Java EE library deployment descriptors (in the order in which they are referenced, one at a time) and then merges the `weblogic-application.xml` file of the referencing Enterprise application on top of the library descriptor files.

As a result of the way the descriptor files are merged, the elements in the descriptors of the shared Java EE libraries referenced first in the `weblogic-application.xml` file have precedence over the ones listed last. The elements of the Enterprise application's descriptor itself have precedence over all elements in the library descriptors.

For example, assume that an Enterprise application called `myApp` references two shared Java EE libraries (themselves packaged as Enterprise applications): `myLibA` and `myLibB`, in that order. Both the `myApp` and `myLibA` applications include an EJB module called `myEJB`, and both the `myLibA` and `myLibB` applications include an EJB module called `myOtherEJB`.

Further assume that once the `myApp` application is deployed, a client invokes, via the `myApp` application, the `myEJB` module. In this case, WebLogic Server actually invokes the EJB in the `myApp` application (rather than the one in `myLibA`) because modules in the *referencing* application have higher precedence over modules in the *referenced* applications. If a client invokes the `myOtherEJB` EJB, then WebLogic Server invokes the one in `myLibA`, because the library is referenced first in the `weblogic-application.xml` file of `myApp`, and thus has precedence over the EJB with the same name in the `myLibB` application.

Best Practices for Using Shared Java EE Libraries

Keep in mind these best practices when developing shared Java EE libraries and optional packages:

- Use shared Java EE Libraries when you want to share one or more Java EE modules (EJBs, Web Applications, Enterprise Applications, or plain Java classes) with multiple Enterprise Applications.

- If you need to deploy a standalone Java EE module, such as an EJB JAR file, as a shared Java EE library, package the module within an Enterprise Application. Doing so avoids potential URI conflicts, because the library URI of a standalone module is derived from the deployment name.

- If you choose to deploy a shared Java EE library as a standalone Java EE module, always specify a known deployment name during deployment and use that name as the URI in referencing applications.

- Use optional packages when multiple Java EE archive files need to share a set of Java classes.

- If you have a set of classes that must be available to applications in an entire domain, and you do not frequently update those classes (for example, if you need to share 3rd party classes in a domain), use the domain `/lib` subdirectory rather than using shared Java EE libraries or optional packages. Classes in the `/lib` subdirectory are added to the system classpath at server start-up time.

- Always specify a specification version and implementation version, even if you do not intend to enforce version requirements with dependent applications. Specifying versions for shared Java EE libraries enables you to deploy multiple versions of the shared files for testing.

- Always specify an `Extension-Name` value for each shared Java EE library. If you do not specify an extension name, one is derived from the deployment name of the library. Default deployment names are different for archive and exploded archive deployments, and they can be set to arbitrary values in the deployment command

- When developing a Web Application for deployment as a shared Java EE library, use a unique context root. If the context root conflicts with the context root in a dependent Java EE application, use the `context-root` element in the EAR's `weblogic-application.xml` deployment descriptor to override the library's context root.

- Package shared Java EE libraries as archive files for delivery to Administrators or deployers in your organization. Deploy libraries from exploded archive directories during development to allow for easy updates and repeated redeployments.

- Deploy shared Java EE libraries to all WebLogic Server instances on which you want to deploy dependent applications and archives. If a library is not registered with a server instance on which you want to deploy a referencing application, deployment of the referencing application fails.

Programming Application Life Cycle Events

The following sections describe how to create applications that respond to WebLogic Server application life cycle events:

- "Understanding Application Life Cycle Events" on page 10-1

- "Registering Events in weblogic-application.xml" on page 10-2

- "Programming Basic Life Cycle Listener Functionality" on page 10-3

- "Examples of Configuring Life Cycle Events with and without the URI Parameter" on page 10-6

- "Understanding Application Life Cycle Event Behavior During Re-deployment" on page 10-7

- "Programming Application Version Life Cycle Events" on page 10-8

WARNING: Application-scoped startup and shutdown classes have been deprecated as of release 9.0 of WebLogic Server. The information in this chapter about startup and shutdown classes is provided only for backwards compatibility. Instead, you should use life cycle listener events in your applications.

Understanding Application Life Cycle Events

Application life cycle listener events provide handles on which developers can control behavior during deployment, undeployment, and redeployment. This section discusses how you can use the application life cycle listener events.

Four application life cycle events are provided with WebLogic Server, which can be used to extend listener, shutdown, and startup classes. These include:

- Listeners—attachable to any event. Possible methods for Listeners are:

  ```
  public void preStart(ApplicationLifecycleEvent evt) {}
  ```

  – The preStart event is the beginning of the prepare phase, or the start of the application deployment process.)

  ```
  public void postStart(ApplicationLifecycleEvent evt) {}
  ```

  – The postStart event is the end of the activate phase, or the end of the application deployment process. The application is deployed.

  ```
  public void preStop(ApplicationLifecycleEvent evt) {}
  ```

  – The preStop event is the beginning of the deactivate phase, or the start of the application removal or undeployment process.

  ```
  public void postStop(ApplicationLifecycleEvent evt) {}
  ```

  – The postStop event is the end of the remove phase, or the end of the application removal or undeployment process.

- Shutdown classes only get postStop events.

  WARNING: Application-scoped shutdown classes have been deprecated as of release 9.0 of WebLogic Server. Use life cycle listeners instead.

- Startup classes only get preStart events.

  WARNING: Application-scoped shutdown classes have been deprecated as of release 9.0 of WebLogic Server. Use life cycle listeners instead.

Note: For Startup and Shutdown classes, you only implement a `main{}` method. If you implement any of the methods provided for Listeners, they are ignored.

No `remove{}` method is provided in the `ApplicationLifecycleListener`, because the events are only fired at startup time during deployment (prestart and poststart) and shutdown during undeployment (prestop and poststop).

Registering Events in weblogic-application.xml

In order to use these events, you must register them in the `weblogic-application.xml` deployment descriptor. See "Application Deployment Descriptor Elements." Define the following elements:

- `listener`—Used to register user defined application life cycle listeners. These are classes that extend the abstract base class `weblogic.application.ApplicationLifecycleListener`.

- `shutdown`—Used to register user-defined shutdown classes.

- `startup`—Used to register user-defined startup classes.

Programming Basic Life Cycle Listener Functionality

You create a listener by extending the abstract class (provided with WebLogic Server) `weblogic.application.ApplicationLifecycleListener`. The container then searches for your listener.

You override the following methods provided in the WebLogic Server `ApplicationLifecycleListener` abstract class to extend your application and add any required functionality:

- preStart{}

- postStart{}

- preStop{}

- postStop{}

Listing 10-1 illustrates how you override the `ApplicationLifecycleListener`. In this example, the public class `MyListener` extends `ApplicationLifecycleListener`.

Listing 10-1 MyListener

```
import weblogic.application.ApplicationLifecycleListener;

import weblogic.application.ApplicationLifecycleEvent;

public class MyListener extends ApplicationLifecycleListener {

  public void preStart(ApplicationLifecycleEvent evt) {

    System.out.println

    ("MyListener(preStart) -- we should always see you..");

  } // preStart

  public void postStart(ApplicationLifecycleEvent evt) {
```

```
      System.out.println

      ("MyListener(postStart) -- we should always see you..");

   } // postStart

  public void preStop(ApplicationLifecycleEvent evt) {

      System.out.println

      ("MyListener(preStop) -- we should always see you..");

   } // preStop

  public void postStop(ApplicationLifecycleEvent evt) {

      System.out.println

      ("MyListener(postStop) -- we should always see you..");

   } // postStop

  public static void main(String[] args) {

      System.out.println

      ("MyListener(main): in main .. we should never see you..");

   } // main

}
```

Listing 10-2 illustrates how you implement the shutdown class. The shutdown class is attachable to preStop and postStop events. In this example, the public class `MyShutdown` does not extend `ApplicationLifecycleListener` because a shutdown class declared in the `weblogic-application.xml` deployment descriptor does not need to depend on any WebLogic Server-specific interfaces.

Listing 10-2 MyShutdown

```
import weblogic.application.ApplicationLifecycleListener;

import weblogic.application.ApplicationLifecycleEvent;

public class MyShutdown {

   public static void main(String[] args) {
```

```
    System.out.println

    ("MyShutdown(main): in main .. should be for post-stop");

  } // main

}
```

Listing 10-3 illustrates how you implement the startup class. The startup class is attachable to preStart and postStart events. In this example, the public class MyStartup does not extend ApplicationLifecycleListener because a startup class declared in the weblogic-application.xml deployment descriptor does not need to depend on any WebLogic Server-specific interfaces.

Listing 10-3   MyStartup

```
import weblogic.application.ApplicationLifecycleListener;

import weblogic.application.ApplicationLifecycleEvent;

public class MyStartup {

  public static void main(String[] args) {

    System.out.println

    ("MyStartup(main): in main .. should be for pre-start");

  } // main

}
```

Configuring a Role-Based Application Life Cycle Listener

You can configure an application life cycle event with role-based capability where a user identity can be specified to startup and shutdown events using the run-as-principal-name element. However, if the run-as-principal-name identity defined for the application life cycle listener is an administrator, the application deployer must have administrator privileges; otherwise, deployment will fail.

1. Follow the basic programming steps outlined in "Programming Basic Life Cycle Listener Functionality" on page 10-3.

2. Within the `listener` element add the `run-as-principal-name` element to specify the user who has privileges to startup and/or shutdown the event. For example:

```
<listener>
  <listener-class>myApp.MySessionAttributeListenerClass</listener-class
>
  <run-as-principal-name>javajoe</run-as-principal-name>
</listener>
```

The identity specified here should be a valid user name in the system. If `run-as-principal-name` is not specified, the deployment initiator user identity will be used as the `run-as` identity for the execution of the application life cycle listener.

Examples of Configuring Life Cycle Events with and without the URI Parameter

The following examples illustrate how you configure application life cycle events in the `weblogic-application.xml` deployment descriptor file. The URI parameter is not required. You can place classes anywhere in the application `$CLASSPATH`. However, you must ensure that the class locations are defined in the `$CLASSPATH`. You can place listeners in `APP-INF/classes` or `APP-INF/lib`, if these directories are present in the EAR. In this case, they are automatically included in the `$CLASSPATH`.

The following example illustrates how you configure application life cycle events using the URI parameter. In this case, the archive `foo.jar` contains the classes and exists at the top level of the EAR file. For example: `myEar/foo.jar`

Listing 10-4   Configuring Application Life Cycle Events Using the URI Parameter

```
<listener>

    <listener-class>MyListener</listener-class>

    <listener-uri>foo.jar</listener-uri>

</listener>

<startup>

    <startup-class>MyStartup</startup-class>
```

```
    <startup-uri>foo.jar</startup-uri>

</startup>

<shutdown>

    <shutdown-class>MyShutdown</shutdown-class>

    <shutdown-uri>foo.jar</shutdown-uri>

</shutdown>
```

The following example illustrates how you configure application life cycle events without using the URI parameter.

Listing 10-5   Configuring Application Life Cycle Events without Using the URI Parameter

```
<listener>

    <listener-class>MyListener</listener-class>

</listener>

<startup>

    <startup-class>MyStartup</startup-class>

</startup>

<shutdown>

    <shutdown-class>MyShutdown</shutdown-class>

</shutdown>
```

Understanding Application Life Cycle Event Behavior During Re-deployment

Application life cycle events are only triggered if a full re-deployment of the application occurs. During a full re-deployment of the application—provided the application life cycle events have

been registered—the application life cycle first commences the shutdown sequence, next re-initializes its classes, and then performs the startup sequence.

For example, if your listener is registered for the full application life cycle set of events (preStart, postStart, preStop, postStop), during a full re-deployment, you see the following sequence of events:

1. `preStop{}`

2. `postStop{}`

3. Initialization takes place. (Unless you have set debug flags, you do not see the initialization.)

4. `preStart{}`

5. `postStart{}`

Programming Application Version Life Cycle Events

The following sections describe how to create applications that respond to WebLogic Server application version life cycle events:

- "Understanding Application Version Life Cycle Event Behavior" on page 10-8

- "Types of Application Version Life Cycle Events" on page 10-9

- "Example of Production Deployment Sequence When Using Application Version Life Cycle Events" on page 10-9

Understanding Application Version Life Cycle Event Behavior

WebLogic Server provides application version life cycle event notifications by allowing you to extend the `ApplicationVersionLifecycleListener` class and specify a life cycle listener in `weblogic-application.xml`. See Application Deployment Descriptor Elements and "Examples of Configuring Life Cycle Events with and without the URI Parameter" on page 10-6.

Application version life cycle events are invoked:

- For both static and dynamic deployments.

- Using either anonymous ID or using user identity.

- Only if the current application is versioned; otherwise, version life cycle events are ignored.

- for all application versions, including the version that registers the listener. Use the `ApplicationVersionLifecycleEvent.isOwnVersion` method to determine if an event belongs to a particular version. See the `ApplicationVersionLifecycleEvent` class for more information on types of version life cycle events.

Types of Application Version Life Cycle Events

Four application version life cycle events are provided with WebLogic Server:

- `public void preDeploy(ApplicationVersionLifecycleEvent evt)`

  – The `preDeloy` event is invoked when an application version deploy or redeploy operation is initiated.

- `public void postDeploy(ApplicationVersionLifecycleEvent evt)`

  – The `postDeloy` event is invoked when an application version is deployed or redeployed successfully.

- `public void preUndeploy(ApplicationVersionLifecycleEvent evt)`

  – The `preUndeloy` event is invoked when an application version undeploy operation is initiated

- `public void postDelete(ApplicationVersionLifecycleEvent evt)`

  – The `postDelete` event is invoked when an application version is deleted.

  Note: A `postDelete` event is only be fired after the entire application version is completely removed. It does not include a partial undeploy, such as undeploying a module or from a subset of targets.

Example of Production Deployment Sequence When Using Application Version Life Cycle Events

The following table provides an example of a deployment (V1), production redeployment (V2), and an undeploy (V2).

Table 10-1 Sequence of Deployment Actions and Application Version Life Cycle Events

Deployment action	Time	Version V1	Version V2
Deployment of Version V1	T0	`preDeploy(V1)` invoked.	
	T1	Deployment starts.	
	T2	Application life cycle listeners for V1 are registered.	
	T3	V1 is active version, Deployment is complete.	
	T4	`postDeploy(V1)` invoked.	
	T5	Application Listeners gets `postDeploy(V1)`.	

Table 10-1 Sequence of Deployment Actions and Application Version Life Cycle Events

Deployment action	Time	Version V1	Version V2
Production Redeployment of Version V2	T6		`preDeploy(V2)` invoked.
	T7	Application version listener receives `preDeploy(V1)`.	
	T8		Deployment starts.
	T9		Application life cycle listeners for V2 are registered.
	T10	If deploy(V2) succeeds, V1 ceases to be active version.	If deploy(V2) succeeds, V2 replaces V1 as active version. Deployment is complete.
	T11		`postDeploy(V2)` invoked. **Note:** This event occurs even if the deployment fails.
	T12	Application version listener gets `postDeploy(V2)`. If deploy(V2) fails, V1 remains active.	
	T13		Application listeners gets `postDeploy(V2)`.
	T14	If deploy(V2) succeeds, V1 begins retirement.	
	T15	Application listeners for V1 are unregistered.	
	T16	V1 is retired.	

Table 10-1 Sequence of Deployment Actions and Application Version Life Cycle Events

Deployment action	Time	Version V1	Version V2
Undeployment of V2	T17		`preUndeploy(v2)` invoked.
	T18		Application listeners gets `preUndeploy(v2)` invoked.
	T19		Undeployment begins.
	T20		V2 is no longer active version.
	T21		Application version listeners for V2 are unregistered.
	T22		Undeployment is complete.
	T23		If the entire application is undeployed, `postDelete(V2)` is invoked.
			Note: This event occurs even if the undeployment fails.

Programming Context Propagation

The following sections describe how to use the context propagation APIs in your applications:

Understanding Context Propagation

Context propagation allows programmers to associate information with an application which is then carried along with every request. Furthermore, downstream components can add or modify this information so that it can be carried back to the originator. Context propagation is also known as *work areas*, *work contexts*, or *application transactions*.

Common use-cases for context propagation are any type of application in which information needs to be carried outside the application, rather than the information being an integral part of the application. Examples of these use cases include diagnostics monitoring, application transactions, and application load-balancing. Keeping this sort of information outside of the application keeps the application itself clean with no extraneous API usage and also allows the addition of information to read-only components, such as 3rd party components.

Programming context propagation has two parts: first you code the client application to create a `WorkContextMap` and `WorkContext`, and then add user data to the context, and then you code the invoked application itself to get and possibly use this data. The invoked application can be of

any type: EJB, Web Service, servlet, JMS topic or queue, and so on. See "Programming Context Propagation: Main Steps" on page 11-3 for details.

The WebLogic context propagation APIs are in the `weblogic.workarea` package. The following table describes the main interfaces and classes.

Table 11-1 Interfaces and classes of the WebLogic Context Propagation API

Interface or Class	Description
`WorkContext Map` Interface	Main context propagation interface used to tag applications with data and propagate that information via application requests. `WorkContextMaps` is part of the client or application's JNDI environment and can be accessed through JNDI by looking up the name `java:comp/WorkContextMap`.
`WorkContext` Interface	Interface used for marshaling and unmarshaling the user data that is passed along with an application. This interface has four implementing classes for marshaling and unmarshaling the following types of data: simple 8-bit ASCII contexts (`AsciiWorkContext`), long contexts (`LongWorkContext`), Serializable context (`SerializableWorkContext`), and String contexts (`StringWorkContext`). `WorkContext` has one subinterface, `PrimitiveWorkContext`, used to specifically marshal and unmarshal a single primitive data item.
`WorkContext Output/Input` Interfaces	Interfaces representing primitive streams used for marshaling and unmarshaling, respectively, `WorkContext` implementations.
`Propagation Mode` Interface	Defines the propagation properties of `WorkContexts`. Specifies whether the WorkContext is propagated locally, across threads, across RMI invocations, across JMS queues and topics, or across SOAP messages. If not specified, default is to propagate data across remote and local calls in the same thread.
`PrimitiveContextFactory` Class	Convenience class for creating `WorkContexts` that contain only primitive data.

For the complete API documentation about context propagation, see the `weblogic.workarea` Javadocs.

Programming Context Propagation: Main Steps

The following procedure describes the high-level steps to use context propagation in your application. It is assumed in the procedure that you have already set up your iterative development environment and have an existing client and application that you want to update to use context propagation by using the `weblogic.workarea` API.

1. Update your client application to create the `WorkContextMap` and `WorkContext` objects and then add user data to the context.

    See "Programming Context Propagation in a Client" on page 11-3.

2. If your client application is standalone (rather than running in a Java EE component deployed to WebLogic Server), ensure that its CLASSPATH includes the Java EE application client, also called the *thin client*.

    See Programming Stand-Alone Clients.

3. Update your application (EJB, Web Service, servlet, and so on) to also create a `WorkContextMap` and then get the context and user data that you added from the client application.

    See "Programming Context Propagation in an Application" on page 11-5.

Programming Context Propagation in a Client

The following sample Java code shows a standalone Java client that invokes a Web Service; the example also shows how to use the `weblogic.workarea.*` context propagation APIs to associate user information with the invoke. The code relevant to context propagation is shown in bold and explained after the example.

For the complete API documentation about context propagation, see the `weblogic.workarea` Javadocs.

Note: See Programming Web Services for WebLogic Server for information on creating Web Services and client applications that invoke them.

```
package examples.workarea.client;

import java.rmi.RemoteException;
```

```
import javax.xml.rpc.ServiceException;
import javax.xml.rpc.Stub;

import javax.naming.InitialContext;
import javax.naming.NamingException;

import weblogic.workarea.WorkContextMap;
import weblogic.workarea.WorkContext;
import weblogic.workarea.PrimitiveContextFactory;
import weblogic.workarea.PropagationMode;
import weblogic.workarea.PropertyReadOnlyException;

/**
 * This is a simple standalone client application that invokes the
 * the <code>sayHello</code> operation of the WorkArea Web service.
 *
 */

public class Main {

  public final static String SESSION_ID= "session_id_key";

  public static void main(String[] args)
      throws ServiceException, RemoteException, NamingException,
PropertyReadOnlyException{

    WorkAreaService service = new WorkAreaService_Impl(args[0] + "?WSDL");
    WorkAreaPortType port = service.getWorkAreaPort();

    WorkContextMap map = (WorkContextMap)new
InitialContext().lookup("java:comp/WorkContextMap");

    WorkContext stringContext = PrimitiveContextFactory.create("A String
Context");

    // Put a string context
    map.put(SESSION_ID, stringContext, PropagationMode.SOAP);

    try {
      String result = null;
      result = port.sayHello("Hi there!");
      System.out.println( "Got result: " + result );
    } catch (RemoteException e) {
      throw e;
    }
  }
}
```

In the preceding example:

- The following code shows how to import the needed `weblogic.workarea.*` classes, interfaces, and exceptions:

```
import weblogic.workarea.WorkContextMap;
import weblogic.workarea.WorkContext;
import weblogic.workarea.PrimitiveContextFactory;
import weblogic.workarea.PropagationMode;
import weblogic.workarea.PropertyReadOnlyException;
```

- The following code shows how to create a `WorkContextMap` by doing a JNDI lookup of the context propagation-specific JNDI name `java:comp/WorkContextMap`:

```
WorkContextMap map = (WorkContextMap)
    new InitialContext().lookup("java:comp/WorkContextMap");
```

- The following code shows how to create a `WorkContext` by using the `PrimitiveContextFactory`. In this example, the `WorkContext` consists of the simple String value `A String Context`. This String value is the user data that is passed to the invoked Web Service.

```
WorkContext stringContext =
    PrimitiveContextFactory.create("A String Context");
```

- Finally, the following code shows how to add the context data, along with the key `SESSION_ID`, to the `WorkContextMap` and associate it with the current thread. The `PropagationMode.SOAP` constant specifies that the propagation happens over SOAP messages; this is because the client is invoking a Web Service.

```
map.put(SESSION_ID, stringContext, PropagationMode.SOAP);
```

Programming Context Propagation in an Application

The following sample Java code shows a simple Java Web Service (JWS) file that implements a Web Service. The JWS file also includes context propagation code to get the user data that is associated with the invoke of the Web Service. The code relevant to context propagation is shown in bold and explained after the example.

For the complete API documentation about context propagation, see the weblogic.workarea Javadocs.

Note: See Programming Web Services for WebLogic Server for information on creating Web Services and client applications that invoke them.

```
package examples.workarea;

import javax.naming.InitialContext;
```

```
// Import the Context Propagation classes

import weblogic.workarea.WorkContextMap;
import weblogic.workarea.WorkContext;

import javax.jws.WebMethod;
import javax.jws.WebService;

import weblogic.jws.WLHttpTransport;

@WebService(name="WorkAreaPortType",
            serviceName="WorkAreaService",
             targetNamespace="http://example.org")

@WLHttpTransport(contextPath="workarea",
                 serviceUri="WorkAreaService",
                  portName="WorkAreaPort")

/**
 * This JWS file forms the basis of simple WebLogic
 * Web Service with a single operation: sayHello
 *
 */

public class WorkAreaImpl {

  public final static String SESSION_ID = "session_id_key";

  @WebMethod()
  public String sayHello(String message) {

   try {

    WorkContextMap map = (WorkContextMap) new
InitialContext().lookup("java:comp/WorkContextMap");

    WorkContext localwc = map.get(SESSION_ID);

    System.out.println("local context: " + localwc);
    System.out.println("sayHello: " + message);

    return "Here is the message: '" + message + "'";

   } catch (Throwable t) {

    return "error";

   }

  }
}
```

In the preceding example:

- The following code shows how to import the needed context propagation APIs; in this case, only the WorkContextMap and WorkContext interfaces are needed:

```
import weblogic.workarea.WorkContextMap;
import weblogic.workarea.WorkContext;
```

- The following code shows how to create a `WorkContextMap` by doing a JNDI lookup of the context propagation-specific JNDI name `java:comp/WorkContextMap`:

```
WorkContextMap map = (WorkContextMap)
    new InitialContext().lookup("java:comp/WorkContextMap");
```

- The following code shows how to get context's user data from the current `WorkContextMap` using a key; in this case, the key is the same one that the client application set when it invoked the Web Service: `SESSION_ID`:

```
WorkContext localwc = map.get(SESSION_ID);
```

Programming Context Propagation

Programming JavaMail with WebLogic Server

The following sections contains information on additional WebLogic Server programming topics:

Overview of Using JavaMail with WebLogic Server Applications

WebLogic Server includes the JavaMail API version 1.3 reference implementation from Sun Microsystems. Using the JavaMail API, you can add email capabilities to your WebLogic Server applications. JavaMail provides access from Java applications to Internet Message Access Protocol (IMAP)- and Simple Mail Transfer Protocol (SMTP)-capable mail servers on your network or the Internet. It does not provide mail server functionality; you must have access to a mail server to use JavaMail.

Complete documentation for using the JavaMail API is available on the JavaMail page on the Sun Web site. This section describes how you can use JavaMail in the WebLogic Server environment.

The `weblogic.jar` file contains the following JavaMail API packages from Sun:

- `javax.mail`

- `javax.mail.event`
- `javax.mail.internet`
- `javax.mail.search`

The `weblogic.jar` also contains the Java Activation Framework (JAF) package, which JavaMail requires.

The `javax.mail` package includes providers for Internet Message Access protocol (IMAP) and Simple Mail Transfer Protocol (SMTP) mail servers. Sun has a separate POP3 provider for JavaMail, which is not included in `weblogic.jar`. You can download the POP3 provider from Sun and add it to the WebLogic Server classpath if you want to use it.

Understanding JavaMail Configuration Files

JavaMail depends on configuration files that define the mail transport capabilities of the system. The `weblogic.jar` file contains the standard configuration files from Sun, which enable IMAP and SMTP mail servers for JavaMail and define the default message types JavaMail can process.

Unless you want to extend JavaMail to support additional transports, protocols, and message types, you do not have to modify any JavaMail configuration files. If you do want to extend JavaMail, download JavaMail from Sun and follow Sun's instructions for adding your extensions. Then add your extended JavaMail package in the WebLogic Server classpath *in front of* `weblogic.jar`.

Configuring JavaMail for WebLogic Server

To configure JavaMail for use in WebLogic Server, you create a mail session in the WebLogic Server Administration Console. This allows server-side modules and applications to access JavaMail services with JNDI, using Session properties you preconfigure for them. For example, by creating a mail session, you can designate the mail hosts, transport and store protocols, and the default mail user in the Administration Console so that modules that use JavaMail do not have to set these properties. Applications that are heavy email users benefit because the mail session creates a single `javax.mail.Session` object and makes it available via JNDI to any module that needs it.

For information on using the Administration Console to create a mail session, see Configure access to JavaMail in the *Administration Console Online Help*.

You can override any properties set in the mail session in your code by creating a `java.util.Properties` object containing the properties you want to override. See "Sending

Then, after you look up the mail session object in JNDI, call the `Session.getInstance()` method with your `Properties` object to get a customized Session.

Sending Messages with JavaMail

Here are the steps to send a message with JavaMail from within a WebLogic Server module:

1. Import the JNDI (naming), JavaBean Activation, and JavaMail packages. You will also need to import `java.util.Properties`:

```
import java.util.*;
import javax.activation.*;
import javax.mail.*;
import javax.mail.internet.*;
import javax.naming.*;
```

2. Look up the Mail Session in JNDI:

```
InitialContext ic = new InitialContext();
Session session = (Session) ic.lookup("myMailSession");
```

3. If you need to override the properties you set for the Session in the Administration Console, create a `java.util.Properties` object and add the properties you want to override. Then call `getInstance()` to get a new Session object with the new properties.

```
Properties props = new Properties();
props.put("mail.transport.protocol", "smtp");
props.put("mail.smtp.host", "mailhost");
// use mail address from HTML form for from address
props.put("mail.from", emailAddress);
Session session2 = session.getInstance(props);
```

4. Construct a `MimeMessage`. In the following example, *to*, *subject*, and *messageTxt* are String variables containing input from the user.

```
Message msg = new MimeMessage(session2);
msg.setFrom();
msg.setRecipients(Message.RecipientType.TO,
                  InternetAddress.parse(to, false));
msg.setSubject(subject);
msg.setSentDate(new Date());
// Content is stored in a MIME multi-part message
// with one body part
MimeBodyPart mbp = new MimeBodyPart();
mbp.setText(messageTxt);
```

```
Multipart mp = new MimeMultipart();
mp.addBodyPart(mbp);
msg.setContent(mp);
```

5. Send the message.

```
Transport.send(msg);
```

The JNDI lookup can throw a `NamingException` on failure. JavaMail can throw a `MessagingException` if there are problems locating transport classes or if communications with the mail host fails. Be sure to put your code in a try block and catch these exceptions.

Reading Messages with JavaMail

The JavaMail API allows you to connect to a message store, which could be an IMAP server or POP3 server. Messages are stored in folders. With IMAP, message folders are stored on the mail server, including folders that contain incoming messages and folders that contain archived messages. With POP3, the server provides a folder that stores messages as they arrive. When a client connects to a POP3 server, it retrieves the messages and transfers them to a message store on the client.

Folders are hierarchical structures, similar to disk directories. A folder can contain messages or other folders. The default folder is at the top of the structure. The special folder name INBOX refers to the primary folder for the user, and is within the default folder. To read incoming mail, you get the default folder from the store, and then get the INBOX folder from the default folder.

The API provides several options for reading messages, such as reading a specified message number or range of message numbers, or pre-fetching specific parts of messages into the folder's cache. See the JavaMail API for more information.

Here are steps to read incoming messages on a POP3 server from within a WebLogic Server module:

1. Import the JNDI (naming), JavaBean Activation, and JavaMail packages. You will also need to import `java.util.Properties`:

```
import java.util.*;
import javax.activation.*;
import javax.mail.*;
import javax.mail.internet.*;
import javax.naming.*;
```

2. Look up the Mail Session in JNDI:

```
InitialContext ic = new InitialContext();
Session session = (Session) ic.lookup("myMailSession");
```

3. If you need to override the properties you set for the Session in the Administration Console, create a `Properties` object and add the properties you want to override. Then call `getInstance()` to get a new Session object with the new properties:

```
Properties props = new Properties();
props.put("mail.store.protocol", "pop3");
props.put("mail.pop3.host", "mailhost");
Session session2 = session.getInstance(props);
```

4. Get a `Store` object from the Session and call its `connect()` method to connect to the mail server. To authenticate the connection, you need to supply the mailhost, username, and password in the connect method:

```
Store store = session.getStore();
store.connect(mailhost, username, password);
```

5. Get the default folder, then use it to get the INBOX folder:

```
Folder folder = store.getDefaultFolder();
folder = folder.getFolder("INBOX");
```

6. Read the messages in the folder into an array of Messages:

```
Message[] messages = folder.getMessages();
```

7. Operate on messages in the Message array. The Message class has methods that allow you to access the different parts of a message, including headers, flags, and message contents.

Reading messages from an IMAP server is similar to reading messages from a POP3 server. With IMAP, however, the JavaMail API provides methods to create and manipulate folders and transfer messages between them. If you use an IMAP server, you can implement a full-featured, Web-based mail client with much less code than if you use a POP3 server. With POP3, you must provide code to manage a message store via WebLogic Server, possibly using a database or file system to represent folders.

Threading and Clustering Topics

The following sections contain information on additional WebLogic Server programming topics:

- "Using Threads in WebLogic Server" on page 13-1
- "Using the Work Manager API for Lower-Level Threading" on page 13-2
- "Programming Applications for WebLogic Server Clusters" on page 13-3

Using Threads in WebLogic Server

WebLogic Server is a sophisticated, multi-threaded application server and it carefully manages resource allocation, concurrency, and thread synchronization for the modules it hosts. To obtain the greatest advantage from WebLogic Server's architecture, construct your application modules created according to the standard Java EE APIs.

In most cases, avoid application designs that require creating new threads in server-side modules:

- Applications that create their own threads do not scale well. Threads in the JVM are a limited resource that must be allocated thoughtfully. Your applications may break or cause WebLogic Server to thrash when the server load increases. Problems such as deadlocks and thread starvation may not appear until the application is under a heavy load.

- Multithreaded modules are complex and difficult to debug. Interactions between application-generated threads and WebLogic Server threads are especially difficult to anticipate and analyze.

In some situations, creating threads may be appropriate, in spite of these warnings. For example, an application that searches several repositories and returns a combined result set can return

results sooner if the searches are done asynchronously using a new thread for each repository instead of synchronously using the main client thread.

If you must use threads in your application code, create a pool of threads so that you can control the number of threads your application creates. Like a JDBC connection pool, you allocate a given number of threads to a pool, and then obtain an available thread from the pool for your runnable class. If all threads in the pool are in use, wait until one is returned. A thread pool helps avoid performance issues and allows you to optimize the allocation of threads between WebLogic Server execution threads and your application.

Be sure you understand where your threads can deadlock and handle the deadlocks when they occur. Review your design carefully to ensure that your threads do not compromise the security system.

To avoid undesirable interactions with WebLogic Server threads, do not let your threads call into WebLogic Server modules. For example, do not use enterprise beans or servlets from threads that you create. Application threads are best used for independent, isolated tasks, such as conversing with an external service with a TCP/IP connection or, with proper locking, reading or writing to files. A short-lived thread that accomplishes a single purpose and ends (or returns to the thread pool) is less likely to interfere with other threads.

Avoid creating daemon threads in modules that are packaged in applications deployed on WebLogic Server. When you create a daemon thread in an application module such as a Servlet, you will not be able to redeploy the application because the daemon thread created in the original deployment will remain running.

Be sure to test multithreaded code under increasingly heavy loads, adding clients even to the point of failure. Observe the application performance and WebLogic Server behavior and then add checks to prevent failures from occurring in production.

Using the Work Manager API for Lower-Level Threading

The Work Manager provides a simple API for concurrent execution of work items. This enables Java EE-based applications (including Servlets and EJBs) to schedule work items for concurrent execution, which will provide greater throughput and increased response time. After an application submits work items to a Work Manager for concurrent execution, the application can gather the results. The Work Manager provides common "join" operations, such as waiting for any or all work items to complete. The Work Manager for Application Servers specification provides an application-server-supported alternative to using lower-level threading APIs, which are inappropriate for use in managed environments such as Servlets and EJBs, as well as being too difficult to use for most applications.

For more information, see Using Work Managers to Optimize Scheduled Work.

Programming Applications for WebLogic Server Clusters

JSPs and Servlets that will be deployed to a WebLogic Server cluster must observe certain requirements for preserving session data. See "Requirements for HTTP Session State Replication" in *Using WebLogic Server Clusters* for more information.

EJBs deployed in a WebLogic Server cluster have certain restrictions based on EJB type. See "Understanding WebLogic Enterprise JavaBeans" in *Programming WebLogic Enterprise JavaBeans* for information about the capabilities of different EJB types in a cluster. EJBs can be deployed to a cluster by setting clustering properties in the EJB deployment descriptor.

If you are developing either EJBs or custom RMI objects for deployment in a cluster, also refer to "Using WebLogic JNDI in a Clustered Environment" in *Programming WebLogic JNDI* to understand the implications of binding clustered objects in the JNDI tree.

Enterprise Application Deployment Descriptor Elements

The following sections describe Enterprise application deployment descriptors: `application.xml` (a Java EE standard deployment descriptor) and `weblogic-application.xml` (a WebLogic-specific application deployment descriptor).

With Java EE annotations, the standard `application.xml` deployment descriptor is optional. Annotations simplify the application development process by allowing developers to specify within the Java class itself how the application component behaves in the container, requests for dependency injection, and so on. Annotations are an alternative to deployment descriptors that were required by older versions of Enterprise applications (J2EE 1.4 and earlier). See Chapter 7, "Using Java EE Annotations and Dependency Injection."

The `weblogic-application.xml` file is also optional if you are not using any WebLogic Server extensions.

- "weblogic-application.xml Deployment Descriptor Elements" on page A-1

- "weblogic-application.xml Schema" on page A-45

- "application.xml Schema" on page A-45

weblogic-application.xml Deployment Descriptor Elements

The following sections describe the many of the individual elements that are defined in the weblogic-application.xml Schema. The `weblogic-application.xml` file is the WebLogic Server-specific deployment descriptor extension for the `application.xml` deployment

descriptor from Sun Microsystems. This is where you configure features such as shared Java EE libraries referenced in the application and EJB caching.

The file is located in the `META-INF` subdirectory of the application archive. The following sections describe elements that can appear in the file.

weblogic-application

The `weblogic-application` element is the root element of the application deployment descriptor.

The following table describes the elements you can define within a `weblogic-application` element.

Table A-1 weblogic-application Elements

Element	Required?	Maximum Number In File	Description
`<ejb>`	Optional	1	Contains information that is specific to the EJB modules that are part of a WebLogic application. Currently, one can use the `ejb` element to specify one or more application level caches that can be used by the application's entity beans. For more information on the elements you can define within the `ejb` element, refer to "ejb" on page A-11.
`<xml>`	Optional	1	Contains information about parsers and entity mappings for XML processing that is specific to this application. For more information on the elements you can define within the `xml` element, refer to "xml" on page A-16.
`<jdbc-connection-pool>`	Optional	Unbounded	Zero or more. Specifies an application-scoped JDBC connection pool. For more information on the elements you can define within the `jdbc-connection-pool` element, refer to "jdbc-connection-pool" on page A-18.
`<security>`	Optional	1	Specifies security information for the application. For more information on the elements you can define within the `security` element, refer to "security" on page A-33.

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<applicatio n-param>`	Optional	Unbounde d	Zero or more. Used to specify un-typed parameters that affect the behavior of container instances related to the application. The parameters listed here are currently supported. Also, these parameters in `weblogic-application.xml` can determine the default encoding to be used for requests and for responses.

- `webapp.encoding.default`—Can be set to a string representing an encoding supported by the JDK. If set, this defines the default encoding used to process servlet requests and servlet responses. This setting is ignored if `webapp.encoding.usevmdefault` is set to `true`. This value is also overridden for request streams by the `input-charset` element of `weblogic.xml`.

- `webapp.encoding.usevmdefault`—Can be set to `true` or `false`. If `true`, the system property `file.encoding` is used to define the default encoding.

The following parameter is used to affect the behavior of Web applications that are contained in this application.

- `webapp.getrealpath.accept_context_pa th`—This is a compatibility switch that may be set to `true` or `false`. If set to `true`, the context path of Web applications is allowed in calls to the servlet API `getRealPath`.

Example:

`<application-param>`

`<param-name>webapp.encoding.default </param-name>`

`<param-value>UTF8</param-value>`

`</application-param>`

For more information on the elements you can define within the `application-param` element, refer to .

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<classloader-structure>`	Optional	Unbounded	A `classloader-structure` element allows you to define the organization of classloaders for this application. The declaration represents a tree structure that represents the classloader hierarchy and associates specific modules with particular nodes. A module's classes are loaded by the classloader that its associated with this element.

Example:

```
<classloader-structure>
     <module-ref>
          <module-uri>ejb1.jar</module-uri>
     </module-ref>

</classloader-structure>
```

```
<classloader-structure>
     <module-ref>
          <module-uri>ejb2.jar</module-uri>
     </module-ref>

</classloader-structure>
```

For more information on the elements you can define within the `classloader-structure` element, refer to "classloader-structure" on page A-34.

Element	Required?	Maximum Number In File	Description
`<listener>`	Optional	Unbounded	Zero or more. Used to register user defined application lifecycle listeners. These are classes that extend the abstract base class `weblogic.application.ApplicationLifecycleListener`.

For more information on the elements you can define within the `listener` element, refer to "listener" on page A-34.

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<startup>`	Optional	Unbounded	Zero or more. Used to register user-defined startup classes.
			For more information on the elements you can define within the startup element, refer to "startup" on page A-35.
			Note: Application-scoped startup and shutdown classes have been deprecated as of release 9.0 of WebLogic Server. Instead, you should use lifecycle listener events in your applications. For details, see Chapter 10, "Programming Application Life Cycle Events."
`<shutdown>`	Optional	Unbounded	Zero or more. Used to register user defined shutdown classes.
			For more information on the elements you can define within the shutdown element, refer to "shutdown" on page A-36.
			Note: Application-scoped startup and shutdown classes have been deprecated as of release 9.0 of WebLogic Server. Instead, you should use lifecycle listener events in your applications. For details, see Chapter 10, "Programming Application Life Cycle Events."

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<module>`	Optional	Unbounded	Represents a single WebLogic application module, such as a JMS or JDBC module.
			This element has the following child elements:
			• `name`—The name of the module.
			• `type`—The type of module. Valid values are JMS, JDBC, or Interception.
			• `path`—The path of the XML file that fully describes the module, relative to the root of the Enterprise application.
			The following example shows how to specify a JMS module called `Workflows`, fully described by the XML file `jms/Workflows-jms.xml`:
			``` <module> <name>Workflows</name> <type>JMS</type> <path>jms/Workflows-jms.xml</path> </module> ```
`<library-ref>`	Optional	Unbounded	A reference to a shared Java EE library.
			For more information on the elements you can define within the `library` element, refer to "library-ref" on page A-43.
`<fair-share-request>`	Optional	Unbounded	Specifies a fair share request class, which is a type of Work Manager request class. In particular, a fair share request class specifies the average percentage of thread-use time required to process requests.
			The `<fair-share-request>` element can take the following child elements:
			• `name`—The name of the fair share request class.
			• `fair-share`—An integer representing the average percentage of thread-use time.
			See Using Work Managers to Optimize Scheduled Work.

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<response-time-request>`	Optional	Unbounded	Specifies a response time request class, which is a a type of Work manager class. In particular, a response time request class specifies a response time goal in milliseconds. The `<response-time-request>` element can take the following child elements: • `name`—The name of the response time request class. • `goal-ms`—The integer response time goal. See Using Work Managers to Optimize Scheduled Work.
`<context-request>`	Optional	Unbounded	Specifies a context request class, which is a a type of Work manager class. In particular, a context request class assigns request classes to requests based on context information, such as the current user or the current user's group. The `<context-request>` element can take the following child elements: • `name`—The name of the context request class. • `context-case`—An element that describes the context. The `<context-case>` element can itself take the following child elements: • `user-name` or `group-name`—The user or group to which the context applies. • `request-class-name`—The name of the request class. See Using Work Managers to Optimize Scheduled Work.

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<max-thread s-constrain t>`	Optional	Unbounde d	Specifies a `max-threads-constraint` Work Manager constraint. A Work Manager constraint defines minimum and maximum numbers of threads allocated to execute requests and the total number of requests that can be queued or executing before WebLogic Server begins rejecting requests.
			The max-threads constraint limits the number of concurrent threads executing requests from the constrained work set.
			The `<max-threads-constraint>` element can take the following child elements:
			• `name`—The name of the max-thread-constaint constraint.
			• Either `count` or `pool-name`—The integer maximum number of concurrent threads, or the name of a connection pool which determines the maximum.
			See Using Work Managers to Optimize Scheduled Work.

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<min-thread s-constrain t>`	Optional	Unbounde d	Specifies a `min-threads-constraint` Work Manager constraint. A Work Manager constraint defines minimum and maximum numbers of threads allocated to execute requests and the total number of requests that can be queued or executing before WebLogic Server begins rejecting requests. The min-threads constraint guarantees a number of threads the server will allocate to affected requests to avoid deadlocks. The `<min-threads-constraint>` element can take the following child elements: • `name`—The name of the min-thread-constraint constraint. • `count`—The integer minimum number of threads. See Using Work Managers to Optimize Scheduled Work.
`<capacity>`	Optional	Unbounde d	Specifies a `capacity` Work Manager constraint. A Work Manager constraint defines minimum and maximum numbers of threads allocated to execute requests and the total number of requests that can be queued or executing before WebLogic Server begins rejecting requests. The capacity constraint causes the server to reject requests only when it has reached its capacity. The `<capacity>` element can take the following child elements: • `name`—The name of the capacity constraint. • `count`—The integer thread capacity. See Using Work Managers to Optimize Scheduled Work.

Table A-1 weblogic-application Elements (Continued)

Element	Required?	Maximum Number In File	Description
`<work-manag er>`	Optional	Unbounde d	Specifies the Work Manager that is associated with the application.
			For more information on the elements you can define within the `work-manager` element, refer to "work-manager" on page A-36.
			See Using Work Managers to Optimize Scheduled Work for detailed information on Work Managers.
`<applicatio n-admin-mod e-trigger>`	Optional	Unbounde d	Specifies the number of stuck threads needed to bring the application into administration mode.
			You can specify the following child elements:
			• `max-stuck-thread-time`—The maximum amount of time, in seconds, that a thread should remain stuck.
			• `stuck-thread-count`—Number of stuck threads that triggers the stuck thread work manager.
`<session-de scriptor>`	Optional	Unbounde d	Specifies a list of configuration parameters for servlet sessions.
			For more information on the elements you can define within the `<session-descriptor>` element, refer to "session-descriptor" on page A-39.
`<fast-swap>`	Optional	1	Specifies whether FastSwap deployment is used minimize redeployment since Java classes are redefined in-place without reloading the ClassLoader.
			For more information, see Using FastSwap Deployment to Minimize Redeployment in *Deploying Applications to WebLogic Server*.

ejb

The following table describes the elements you can define within an `ejb` element.

Table A-2 ejb Elements

Element	Required?	Maximum Number in File	Description
`<entity-cache>`	Optional	Unbounded	Zero or more. The `entity-cache` element is used to define a named application level cache that is used to cache entity EJB instances at runtime. Individual entity beans refer to the application-level cache that they must use, referring to the cache name. There is no restriction on the number of different entity beans that may reference an individual cache.
			To use application-level caching, you must specify the cache using the `<entity-cache-ref>` element of the `weblogic-ejb-jar.xml` descriptor. Two default caches named `ExclusiveCache` and `MultiVersionCache` are used for this purpose. An application may explicitly define these default caches to specify non-default values for their settings. Note that the caching-strategy cannot be changed for the default caches. By default, a cache uses `max-beans-in-cache` with a value of 1000 to specify its maximum size.
			Example:
			`<entity-cache>`
			`<entity-cache-name>ExclusiveCache</entity-cache-name>`
			`<max-cache-size>`
			`<megabytes>50</megabytes>`
			`</max-cache-size>`
			</entity-cache>
			For more information on the elements you can define within the `entity-cache` element, refer to "entity-cache" on page A-13.

Table A-2 ejb Elements

Element	Required?	Maximum Number in File	Description
`<start-mbds-wi th- application`	Optional	1	Allows you to configure the EJB container to start Message Driven BeanS (MDBS) with the application. If set to true, the container starts MDBS as part of the application. If set to false, the container keeps MDBS in a queue and the server starts them as soon as it has started listening on the ports.

entity-cache

The following table describes the elements you can define within a `entity-cache` element.

Table A-3 entity-cache Elements

Element	Required?	Maximum Number in File	Description
`<entity-cache-name>`	Required	1	Specifies a unique name for an entity bean cache. The name must be unique within an ear file and may not be the empty string. Example: `<entity-cache-name>ExclusiveCache</enti ty-cache-name>`
`<max-beans-in-cache>`	Optional If you specify this element, you cannot also specify `<max-cac he-size>`.	1	Specifies the maximum number of entity beans that are allowed in the cache. If the limit is reached, beans may be passivated. This mechanism does not take into account the actual amount of memory that different entity beans require. This element can be set to a value of 1 or greater. Default Value: `1000`

Table A-3 entity-cache Elements

Element	Required?	Maximum Number in File	Description
`<max-cache-siz e>`	Optional If you specify this element, you cannot also specify `<max-bea ns-in-ca che>`.	1	Used to specify a limit on the size of an entity cache in terms of memory size—expressed either in terms of bytes or megabytes. A bean provider should provide an estimate of the average size of a bean in the `weblogic-ejb-jar.xml` descriptor if the bean uses a cache that specifies its maximum size using the `max-cache-size` element. By default, a bean is assumed to have an average size of 100 bytes. For more information on the elements you can define within the `ejb` element, refer to "max-cache-size" on page A-15.
`<max-queries-i n-cache>`	Optional	1	Specifies the maximum SQL queries that can be present in the entity cache at a given moment.
`<caching-strat egy>`	Optional	1	Specifies the general strategy that the EJB container uses to manage entity bean instances in a particular application level cache. A cache buffers entity bean instances in memory and associates them with their primary key value. The `caching-strategy` element can only have one of the following values: • `Exclusive`—Caches a single bean instance in memory for each primary key value. This unique instance is typically locked using the EJB container's exclusive locking when it is in use, so that only one transaction can use the instance at a time. • `MultiVersion`—Caches multiple bean instances in memory for a given primary key value. Each instance can be used by a different transaction concurrently. Default Value: `MultiVersion` Example: `<caching-strategy>Exclusive</caching-strategy>`

max-cache-size

The following table describes the elements you can define within a `max-cache-size` element.

Table A-4 max-cache-size Elements

Element	Required?	Maximum Number in File	Description
`<bytes>`	You *must* specify either `<bytes>` or `<megabytes>`	1	The size of an entity cache in terms of memory size, expressed in bytes.
`<megabytes>`	You *must* specify either `<bytes>` or `<megabytes>`	1	The size of an entity cache in terms of memory size, expressed in megabytes.

xml

The following table describes the elements you can define within an `xml` element.

Table A-5 xml Elements

Element	Required?	Maximum Number in File	Description
`<parser-factory>`	Optional	1	The parent element used to specify a particular XML parser or transformer for an enterprise application. For more information on the elements you can define within the `parser-factory` element, refer to "parser-factory" on page A-17.
`<entity-mapping>`	Optional	Unbounded	Zero or More. Specifies the entity mapping. This mapping determines the alternative entity URI for a given public or system ID. The default place to look for this entity URI is the `lib/xml/registry` directory. For more information on the elements you can define within the `entity-mapping` element, refer to "entity-mapping" on page A-18.

parser-factory

The following table describes the elements you can define within a `parser-factory` element.

Table A-6 parser-factory Elements

Element	Required?	Maximum Number in File	Description
`<saxparser-factory>`	Optional	1	Allows you to set the SAXParser Factory for the XML parsing required in this application only. This element determines the factory to be used for SAX style parsing. If you do not specify the `saxparser-factory` element setting, the configured SAXParser Factory style in the Server XML Registry is used. Default Value: Server XML Registry setting
`<document-builder-factory>`	Optional	1	Allows you to set the Document Builder Factory for the XML parsing required in this application only. This element determines the factory to be used for DOM style parsing. If you do not specify the `document-builder-factory` element setting, the configured DOM style in the Server XML Registry is used. Default Value: Server XML Registry setting
`<transformer-factory>`	Optional	1	Allows you to set the Transformer Engine for the style sheet processing required in this application only. If you do not specify a value for this element, the value configured in the Server XML Registry is used. Default value: Server XML Registry setting.

entity-mapping

The following table describes the elements you can define within an `entity-mapping` element.

Table A-7 entity-mapping Elements

Element	Required?	Maximum Number in File	Description
`<entity-mapping-name>`	Required	1	Specifies the name for this entity mapping.
`<public-id>`	Optional	1	Specifies the public ID of the mapped entity.
`<system-id>`	Optional	1	Specifies the system ID of the mapped entity.
`<entity-uri>`	Optional	1	Specifies the entity URI for the mapped entity.
`<when-to-cache>`	Optional	1	Legal values are: • cache-on-reference • cache-at-initialization • cache-never The default value is `cache-on-reference`.
`<cache-timeout-interval>`	Optional	1	Specifies the integer value in seconds.

jdbc-connection-pool

Note: The `jdbc-connection-pool` element is deprecated. To define a data source in your Enterprise application, you can package a JDBC module with the application. For more information, see Configuring JDBC Application Modules for Deployment in *Configuring and Managing WebLogic JDBC*.

The following table describes the elements you can define within a `jdbc-connection-pool` element.

Table A-8 jdbc-connection-pool Elements

Element	Required?	Maximum Number in File	Description
`<data-source-jndi-name>`	Required	1	Specifies the JNDI name in the application-specific JNDI tree.
`<connection-factory>`	Required	1	Specifies the connection parameters that define overrides for default connection factory settings. • `user-name`—Optional. The `user-name` element is used to override `UserName` in the `JDBCDataSourceFactoryMBean`. • `url`—Optional. The `url` element is used to override `URL` in the `JDBCDataSourceFactoryMBean`. • `driver-class-name`—Optional. The `driver-class-name` element is used to override `DriverName` in the `JDBCDataSourceFactoryMBean`. • `connection-params`—Zero or more. • `parameter+` (param-value, param-name)—One or more For more information on the elements you can define within the `connection-factory` element, refer to "connection-factory" on page A-20.
`<pool-params>`	Optional	1	Defines parameters that affect the behavior of the pool. For more information on the elements you can define within the `pool-params` element, refer to "pool-params" on page A-21.
`<driver-params>`	Optional	1	Sets behavior on WebLogic Server drivers. For more information on the elements you can define within the `driver-params` element, refer to "driver-params" on page A-29.
`<acl-name>`	Optional	1	DEPRECATED.

connection-factory

The following table describes the elements you can define within a `connection-factory` element.

Table A-9 connection-factory Elements

Element	Required?	Maximum Number in File	Description
`<factory-name>`	Optional	1	Specifies the name of a `JDBCDataSourceFactoryMBean` in the `config.xml` file.

`<connection-pr operties>`	Optional	1	Specifies the connection properties for the connection factory. Elements that can be defined for the `connection-properties` element are:
			
			• `user-name`—Optional. Used to override UserName in the JDBCDataSourceFactoryMBean.
			
			• `password`—Optional. Used to override Password in the JDBCDataSourceFactoryMBean.
			
			• `url`—Optional. Used to override URL in the JDBCDataSourceFactoryMBean.
			
			• `driver-class-name`—Optional. Used to override DriverName in the JDBCDataSourceFactoryMBean
			
			• `connection-params`—Zero or more. Used to set parameters which will be passed to the driver when making a connection. Example:
			
			`<connection-params>`
			`  <parameter>`
			`   <description>Desc of param`
			`   </description>`
			`   <param-name>foo</param-name>`
			`   <param-value>xyz</param-value>`
			`  </parameter>`
			`</connection-params>`

pool-params

The following table describes the elements you can define within a `pool-params` element.

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description
`<size-params>`	Optional	1	Defines parameters that affect the number of connections in the pool.
			• `initial-capacity`—Optional. The `initial-capacity` element defines the number of physical database connections to create when the pool is initialized. The default value is `1`.
			• `max-capacity`—Optional. The `max-capacity` element defines the maximum number of physical database connections that this pool can contain. Note that the JDBC Driver may impose further limits on this value. The default value is `1`.
			• `capacity-increment`—Optional. The `capacity-increment` element defines the increment by which the pool capacity is expanded. When there are no more available physical connections to service requests, the pool creates this number of additional physical database connections and adds them to the pool. The pool ensures that it does not exceed the maximum number of physical connections as set by `max-capacity`. The default value is `1`.
			• `shrinking-enabled`—Optional. The `shrinking-enabled` element indicates whether or not the pool can shrink back to its `initial-capacity` when connections are detected to not be in use.
			• `shrink-period-minutes`—Optional. The `shrink-period-minutes` element defines the number of minutes to wait before shrinking a connection pool that has incrementally increased to meet demand. The `shrinking-enabled` element must be set to `true` for shrinking to take place.
			• `shrink-frequency-seconds`—Optional.
			• `highest-num-waiters`—Optional.
			• `highest-num-unavailable`—Optional.

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description
`<xa-params>`	Optional	1	Defines the parameters for the XA DataSources.

- `debug-level`—Optional. Integer. The `debug-level` element defines the debugging level for XA operations. The default value is `0`.

- `keep-conn-until-tx-complete-enabled`—Optional. Boolean. If you set the `keep-conn-until-tx-complete-enabled` element to `true`, the XA connection pool associates the same XA connection with the distributed transaction until the transaction completes.

- `end-only-once-enabled`—Optional. Boolean. If you set the `end-only-once-enabled` element to `true`, the `XAResource.end()` method is only called once for each pending `XAResource.start()` method.

- `recover-only-once-enabled`—Optional. Boolean. If you set the recover-only-once-enabled element to true, recover is only called one time on a resource.

- `tx-context-on-close-needed`—Optional. Set the `tx-context-on-close-needed` element to `true` if the XA driver requires a distributed transaction context when closing various JDBC objects (for example, result sets, statements, connections, and so on). If set to `true`, the SQL exceptions that are thrown while closing the JDBC objects in no transaction context are swallowed.

- `new-conn-for-commit-enabled`—Optional. Boolean. If you set the `new-conn-for-commit-enabled` element to `true`, a dedicated XA connection is used for commit/rollback processing of a particular distributed transaction.

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description
`<xa-params>` Continued...	Optional	1	• `prepared-statement-cache-size`— **Deprecated**. Optional. Use the prepared-statement-cache-size element to set the size of the prepared statement cache. The size of the cache is a number of prepared statements created from a particular connection and stored in the cache for further use. Setting the size of the prepared statement cache to 0 turns it off.
			Note: `Prepared-statement-cache-size` is deprecated. Use `cache-size` in `driver-params/prepared-statement`. See "driver-params" for more information.
			• `keep-logical-conn-open-on-release`— Optional. Boolean. Set the `keep-logical-conn-open-on-release` element to `true`, to keep the logical JDBC connection open when the physical XA connection is returned to the XA connection pool. The default value is `false`.
			• `local-transaction-supported`—Optional. Boolean. Set the `local-transaction-supported` to `true` if the XA driver supports SQL with no global transaction; otherwise, set it to `false`. The default value is `false`.
			• `resource-health-monitoring-enabled`—Optional. Set the `resource-health-monitoring-enabled` element to `true` to enable JTA resource health monitoring for this connection pool.

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description

`<xa-params>` Continued...	Optional	1	• `xa-set-transaction-timeout`—Optional.
			
			Used in: `xa-params`
			
			Example:
			
			    <xa-set-transaction-timeout>
			     true
			    </xa-set-transaction-timeout>
			
			• `xa-transaction-timeout`—Optional.
			
			When the `xa-set-transaction-timeout` value is set to true, the transaction manager invokes setTransactionTimeout on the resource before calling `XAResource.start`. The Transaction Manager passes the global transaction timeout value. If this attribute is set to a value greater than `0`, then this value is used in place of the global transaction timeout.
			
			Default value: `0`
			
			Used in: `xa-params`
			
			Example:
			
			    <xa-transaction-timeout>
			     30
			    </xa-transaction-timeout>
			
			• `rollback-localtx-upon-connclose`— Optional.
			
			When the `rollback-localtx-upon-connclose` element is true, the connection pool calls `rollback()` on the connection before putting it back in the pool.
			
			Default value: `false`
			
			Used in: `xa-params`
			
			Example:
			
			<rollback-localtx-upon-connclose>
			          true
			</rollback-localtx-upon-connclose>

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description
`<login-delay-s econds>`	Optional	1	Sets the number of seconds to delay before creating each physical database connection. Some database servers cannot handle multiple requests for connections in rapid succession. This property allows you to build in a small delay to let the database server catch up. This delay occurs both during initial pool creation and during the lifetime of the pool whenever a physical database connection is created.
`<leak-profilin g-enabled>`	Optional	1	Enables JDBC connection leak profiling. A connection leak occurs when a connection from the pool is not closed explicitly by calling the `close()` method on that connection. When connection leak profiling is active, the pool stores the stack trace at the time the connection object is allocated from the pool and given to the client. When a connection leak is detected (when the connection object is garbage collected), this stack trace is reported. This element uses extra resources and will likely slowdown connection pool operations, so it is not recommended for production use.

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description
`<connection-ch eck-params>`	Optional	1	• Defines whether, when, and how connections in a pool is checked to make sure they are still alive.
			• `table-name`—Optional. The `table-name` element defines a table in the schema that can be queried.
			• `check-on-reserve-enabled`—Optional. If the check-on-reserve-enabled element is set to true, then the connection will be tested each time before it is handed out to a user.
			• `check-on-release-enabled`—Optional. If the `check-on-release-enabled` element is set to `true`, then the connection will be tested each time a user returns a connection to the pool.
			• `refresh-minutes`—Optional. If the `refresh-minutes` element is defined, a trigger is fired periodically (based on the number of minutes specified). This trigger checks each connection in the pool to make sure it is still valid.
			• `check-on-create-enabled`—Optional. If set to `true`, then the connection will be tested when it is created.
			• `connection-reserve-timeout-seconds`—Optional. Number of seconds after which the call to reserve a connection from the pool will timeout.
			• `connection-creation-retry-frequency-s econds`—Optional. The frequency of retry attempts by the pool to establish connections to the database.
			• `inactive-connection-timeout-seconds`— Optional. The number of seconds of inactivity after which reserved connections will forcibly be released back into the pool.

Table A-10 pool-params Elements

Element	Required?	Maximum Number in File	Description
`<connection-check-params>` Continued...	Optional	1	• `test-frequency-seconds`—Optional. The number of seconds between database connection tests. After every test-frequency-seconds interval, unused database connections are tested using `table-name`. Connections that do not pass the test will be closed and reopened to re-establish a valid physical database connection. If `table-name` is not set, the test will not be performed. • `init-sql`—Optional. Specifies a SQL query that automatically runs when a connection is created.
`<jdbcxa-debug-level>`	Optional	1	This is an internal setting.
`<remove-infected-connections-enabled>`	Optional	1	Controls whether a connection is removed from the pool when the application asks for the underlying vendor connection object. Enabling this attribute has an impact on performance; it essentially disables the pooling of connections (as connections are removed from the pool and replaced with new connections).

driver-params

The following table describes the elements you can define within a `driver-params` element.

Table A-11 driver-params Elements

Element	Required?	Maximum Number in File	Description

`<statement>`	Optional	1	Defines the `driver-params` statement. Contains the following optional element: `profiling-enabled`.
			
			Example:
			
			`<statement>`
			  `<profiling-enabled>true`
			  `</profiling-enabled>`
			`</statement>`

Table A-11 driver-params Elements

Element	Required?	Maximum Number in File	Description
`<prepared-stat ement`	Optional	1	Enables the running of JDBC prepared statement cache profiling. When enabled, prepared statement cache profiles are stored in external storage for further analysis. This is a resource-consuming feature, so it is recommended that you turn it off on a production server. The default value is false.

- `profiling-enabled`—Optional.
- `cache-profiling-threshold`—Optional. The `cache-profiling-threshold` element defines a number of statement requests after which the state of the prepared statement cache is logged. This element minimizes the output volume. This is a resource-consuming feature, so it is recommended that you turn it off on a production server.
- `cache-size`—Optional. The `cache-size` element returns the size of the prepared statement cache. The size of the cache is a number of prepared statements created from a particular connection and stored in the cache for further use.
- `parameter-logging-enabled`—Optional. During SQL roundtrip profiling it is possible to store values of prepared statement parameters. The `parameter-logging-enabled` element enables the storing of statement parameters. This is a resource-consuming feature, so it is recommended that you turn it off on a production server.
- `max-parameter-length`—Optional. During SQL roundtrip profiling it is possible to store values of prepared statement parameters. The `max-parameter-length` element defines maximum length of the string passed as a parameter for JDBC SQL roundtrip profiling. This is a resource-consuming feature, so you should limit the length of data for a parameter to reduce the output volume.
- `cache-type`—Optional.

Table A-11 driver-params Elements

Element	Required?	Maximum Number in File	Description
`<row-prefetch-enabled>`	Optional	1	Specifies whether to enable row prefetching between a client and WebLogic Server for each ResultSet.
			When an external client accesses a database using JDBC through Weblogic Server, row prefetching improves performance by fetching multiple rows from the server to the client in one server access. WebLogic Server ignores this setting and does not use row prefetching when the client and WebLogic Server are in the same JVM
`<row-prefetch-size>`	Optional	1	Specifies the number of result set rows to prefetch for a client.
			The optimal value depends on the particulars of the query. In general, increasing this number increases performance, until a particular value is reached. At that point further increases do not result in any significant increase in performance.
			Note: Typically you will not see any increase in performance after 100 rows. The default value should be adequate for most situations.
			Valid values for this element are between 2 and 65536. The default value is 48.
`<stream-chunk-size>`	Optional	1	Specifies the data chunk size for streaming data types, which are pulled from WebLogic Server to the client as needed.

security

The following table describes the elements you can define within a security element.

Table A-12 security Elements

Element	Required?	Maximum Number in File	Description
`<realm-name>`	Optional	1	Names a security realm to be used by the application. If none is specified, the system default realm is used

`<security-role -assignment>`	Optional	Unbounded	Declares a mapping between an application-wide security role and one or more WebLogic Server principals.
			
			Example:
			
			`<security-role-assignment>`
			`  <role-name>`
			`    PayrollAdmin`
			`  </role-name>`
			`  <principal-name>`
			`    Tanya`
			`  </principal-name>`
			`  <principal-name>`
			`    Fred`
			`  </principal-name>`
			`  <principal-name>`
			`    system`
			`  </principal-name>`
			`</security-role-assignment>`

application-param

The following table describes the elements you can define within a application-param element.

Table A-13 application-param Elements

Element	Required?	Maximum Number in File	Description
`<description>`	Optional	1	Provides a description of the application parameter.
`<param-name>`	Required	1	Defines the name of the application parameter.
`<param-value>`	Required	1	Defines the value of the application parameter.

classloader-structure

The following table describes the elements you can define within a classloader-structure element.

Table A-14 classloader-structure Elements

Element	Required?	Maximum Number in File	Description
`<module-ref>`	Optional	Unbounded	The following list describes the elements you can define within a `module-ref` element: • `module-uri`—Zero or more. Defined within the `module-ref` element.
`<classloader-structure>`	Optional	Unbounded	Allows for arbitrary nesting of classloader structures for an application. However, for this version of WebLogic Server, the depth is restricted to three levels.

listener

The following table describes the elements you can define within a listener element.

Table A-15 listener Elements

Element	Required?	Maximum Number in File	Description
`<listener-class>`	Required	1	Name of the user's implementation of `ApplicationLifecycleListener`.
`<listener-uri>`	Optional	1	A JAR file within the EAR that contains the implementation. If you do not specify the `listener-uri`, it is assumed that the class is visible to the application.
`<run-as-principal-name>`	Optional	1	Specific a user identity to startup and shutdown application lifecycle events. The identity specified here should be a valid user name in the system. If `run-as-principal-name` is not specified, the deployment initiator user identity will be used as the `run-as` identity for the execution of the application lifecycle listener.
			Note: If the `run-as-principal-name` identity defined for the application lifecycle listener is an administrator, the application deployer must have administrator privileges; otherwise, deployment will fail.

startup

The following table describes the elements you can define within a startup element.

WARNING: Application-scoped startup and shutdown classes have been deprecated as of release 9.0 of WebLogic Server. Instead, you should use lifecycle listener events in your applications. For details, see Chapter 10, "Programming Application Life Cycle Events."

Table A-16 startup Elements

Element	Required?	Maximum Number in File	Description
`<startup-clas s>`	Required	1	Defines the name of the class to be run when the application is being deployed.
`<startup-uri>`	Optional	1	Defines a JAR file within the EAR that contains the `startup-class`. If `startup-uri` is not defined, then its assumed that the class is visible to the application.

shutdown

The following table describes the elements you can define within a `shutdown` element.

WARNING: Application-scoped startup and shutdown classes have been deprecated as of release 9.0 of WebLogic Server. Instead, you should use lifecycle listener events in your applications. For details, see Chapter 10, "Programming Application Life Cycle Events."

Table A-17 shutdown Elements

Element	Required Optional	Maximum Number in File	Description
`<shutdown-cla ss>`	Required	1	Defines the name of the class to be run when the application is undeployed.
`<shutdown-uri >`	Optional	1	Defines a JAR file within the EAR that contains the `shutdown-class`. If you do not define the `shutdown-uri` element, it is assumed that the class is visible to the application.

work-manager

The following table describes the elements you can define within a work-manager element.

See Using Work Managers to Optimize Scheduled Work for examples and information on Work Managers.

Table A-18 work-manager Elements

Element	Required?	Maximum Number in File	Description
`<name>`	Required	1	The name of the Work Manager.
`<response-time-request-class>`	Optional	1	See the description of the `<response-time-request>` element in "weblogic-application" on page A-2 for information on this child element of `<work-manager>`. If you specify this element, you cannot also specify `<fair-share-request-class>`, `<context-request-class>`, or `<request-class-name>`.
`<fair-share-request-class>`	Optional	1	See the description of the `<fair-share-request>` element in "weblogic-application" on page A-2 for information on this child element of `<work-manager>`. If you specify this element, you cannot also specify `<response-time-request-class>`, `<context-request-class>`, or `<request-class-name>`.
`<context-request-class>`	Optional	1	See the description of the `<context-request>` element in "weblogic-application" on page A-2 for information on this child element of `<work-manager>`. If you specify this element, you cannot also specify `<fair-share-request-class>`, `<response-time-request-class>`, or `<request-class-name>`.

Table A-18 work-manager Elements

Element	Required?	Maximum Number in File	Description
`<request-class-name>`	Optional	1	The name of the request class.
			If you specify this element, you cannot also specify `<fair-share-request-class>`, `<context-request-class>`, or `<response-time-request-class>`.
`<min-threads-constraint>`	Optional	1	See the description of the `<min-threads-constraint>` element in "weblogic-application" on page A-2 for information on this child element of `<work-manager>`.
			If you specify this element, you cannot also specify `<min-threads-constaint-name>`.
`<min-threads-constraint-name>`	Optional	1	The name of the min-threads constraint.
			If you specify this element, you cannot also specify `<min-threads-constaint>`.
`<max-threads-constraint>`	Optional	1	See the description of the `<max-threads-constraint>` element in "weblogic-application" on page A-2 for information on this child element of `<work-manager>`.
			If you specify this element, you cannot also specify `<max-threads-constaint-name>`.
`<max-threads-constraint-name>`	Optional	1	The name of the max-threads constraint.
			If you specify this element, you cannot also specify `<max-threads-constaint>`.
`<capacity>`	Optional	1	See the description of the `<capacity>` element in "weblogic-application" on page A-2 for information on this child element of `<work-manager>`.
			If you specify this element, you cannot also specify `<capacity-name>`.

Table A-18 work-manager Elements

Element	Required?	Maximum Number in File	Description
`<capacity-name>`	Optional	1	The name of the thread capacity constraint.
			If you specify this element, you cannot also specify `<capacity>`.
`<work-manager-shutdown-trigger>`	Optional	1	Used to specify a Stuck Thread Work Manager component that can shut down the Work Manager in response to stuck threads.
			You can specify the following child elements:
			• `max-stuck-thread-time`—The maximum amount of time, in seconds, that a thread should remain stuck.
			• `stuck-thread-count`—Number of stuck threads that triggers the stuck thread work manager.
			If you specify this element, you cannot also specify `<ignore-stuck-threads>`.
`<ignore-stuck-threads>`	Optional	1	Specifies whether the Work Manager should ignore stuck threads and never shut down even if threads become stuck.
			If you specify this element, you cannot also specify `<work-manager-shutdown-trigger>`.

session-descriptor

The following table describes the elements you can define within a session-descriptor element.

Table A-19 session-descriptor Elements

Element	Required?	Maximum Number in File	Description
`<timeout-secs>`	Optional	1	Specifies the number of seconds after which the session times out. Default value is 3600 seconds.
`<invalidation-interval-secs>`	Optional	1	Specifies the number of seconds of the invalidation trigger interval. Default value is 60 seconds.
`<debug-enabled>`	Optional	1	Specifies whether debugging is enabled for HTTP sessions. Default value is `false`.
`<id-length>`	Optional	1	Specifies the length of the session ID. Default value is 52.
`<tracking-enabled>`	Optional	1	Specifies whether session tracking is enabled between HTTP requests. Default value is `true`.
`<cache-size>`	Optional	1	Specifies the cache size for JDBC and file persistent sessions. Default value is 1028.
`<max-in-memory-sessions>`	Optional	1	Specifies the maximum sessions limit for memory/replicated sessions. Default value is -1, or unlimited.
`<cookies-enabled>`	Optional	1	Specifies the Web application container should set cookies in the response. Default value is `true`.
`<cookie-name>`	Optional	1	Specifies the name of the cookie that tracks sessions. Default name is `JSESSIONID`.

Table A-19 session-descriptor Elements

Element	Required?	Maximum Number in File	Description
`<cookie-path>`	Optional	1	Specifies the session tracking cookie path. Default value is `/`.
`<cookie-domain>`	Optional	1	Specifies the session tracking cookie domain. Default value is `null`.
`<cookie-comment>`	Optional	1	Specifies the session tracking cookie comment. Default value is `null`.
`<cookie-secure>`	Optional	1	Specifies whether the session tracking cookie is marked secure. Default value is `false`.
`<cookie-max-age-secs>`	Optional	1	Specifies that maximum age of the session tracking cookie. Default value is `-1`, or unlimited.
`<persistent-store-type>`	Optional	1	Specifies the type of storage for session persistence. You can specify the following values: • `memory`—Default value. • `replicated`—Requires clustering. • `replicated_if_clustered`—Defaults to `memory` in non-clustered case. • `file` • `jdbc` • `cookie`
`<persistent-store-cookie-name>`	Optional	1	Specifies the name of the cookie that holds the attribute name and values when using `cookie`-based session persistence. Default value is `WLCOOKIE`.

Table A-19 session-descriptor Elements

Element	Required?	Maximum Number in File	Description
`<persistent-store-dir>`	Optional	1	Specifies the name of the directory when using `file`-based session persistence. The directory is relative to the temporary directory defined for the Web application. Default value is `session_db`.
`<persistent-store-pool>`	Optional	1	Specifies the name of the JDBC connection pool when using `jdbc`-based session persistence.
`<persistent-store-table>`	Optional	1	Specifies the name of the database table when using `jdbc`-based session persistence. Default value is `wl_servlet_sessions`.
`<jdbc-column-name-max-inactive-interval>`	Optional	1	Alternative name for the `wl_max_inactive_interval` column name when using `jdbc`-based session persistence. Required for certain databases that do not support long column names
`<jdbc-connection-timeout-secs>`	Optional	1	DEPRECATED
`<url-rewriting-enabled>`	Optional	1	Specifies whether URL rewriting is enabled. Default value is `true`.
`<http-proxy-caching-of-cookies>`	Optional	1	Specifies whether WebLogic Server adds the following HTTP header to the response: `    Cache-control: no-cache=set-cookie` This header specifies that proxy caches should not cache the cookies. Default value is `true`, which means that the header is NOT added. Set this element to `false` if you want the header added to the response.

Table A-19 session-descriptor Elements

Element	Required?	Maximum Number in File	Description
`<encode-sessi on-id-in-quer y-params>`	Optional	1	Specifies whether WebLogic Server should encode the session ID in the path parameters. Default value is `false`.
`<monitoring-a ttribute-name >`	Optional	1	Used to tag runtime information for different sessions. For example, set this element to `username` if you have a `username` attribute that is guaranteed to be unique.
`<sharing-enab led>`	Optional	1	Specifies whether HTTP sessions are shared across multiple Web applications. Default value is `false`.

library-ref

The following table describes the elements you can define within a `library-ref` element.

See Chapter 9, "Creating Shared Java EE Libraries and Optional Packages," for additional information and examples.

Table A-20 library Elements

Element	Required?	Maximum Number in File	Description
`<library-name>`	Required	1	Specifies the name of the referenced shared Java EE library.
`<specification -version>`	Optional	1	Specifies the minimum specification-version required.
`<implementatio n-version>`	Optional	1	Specifies the minimum implementation-version required.

Table A-20 library Elements

Element	Required?	Maximum Number in File	Description
`<exact-match>`	Optional	1	Specifies whether there must be an exact match between the specification and implementation version that is specified and that of the referenced library. Default value is `false`.
`<context-root>`	Optional	1	Specifies the context-root of the references Web Applications shared Java EE library.

library-context-root-override

The following table describes the elements you can define within a library-context-root-override element to override `context-root` elements within a referenced EAR library. See "library-ref" on page A-43.

See Chapter 9, "Creating Shared Java EE Libraries and Optional Packages," for additional information and examples.

Table A-21 library-context-root-override Elements

Element	Required?	Maximum Number in File	Description
`<library-context-root>`	Optional	1	Overrides the `context-root` elements declared in libraries. In the absence of this element, the library's `context-root` is used. Only a referencing application (i.e., a user application) can override the `context-root` elements declared in its libraries.
`<override-value>`	Optional	1	Specifies the value of the `library-context-root-override` element when overriding the context-root elements declared in libraries. In the absence of these elements, the library's `context-root` is used.

weblogic-application.xml Schema

See http://www.bea.com/ns/weblogic/weblogic-application/1.0/weblogic-application.xsd for the XML Schema of the `weblogic-application.xml` deployment descriptor file.

application.xml Schema

For more information about `application.xml` deployment descriptor elements, see the JEE 5 schema available at http://java.sun.com/xml/ns/javaee/application_5.xsd.

wldeploy Ant Task Reference

The following sections describe tools for deploying applications and standalone modules to WebLogic Server:

- "Overview of the wldeploy Ant Task" on page B-1

- "Basic Steps for Using wldeploy" on page B-2

- "Sample build.xml Files for wldeploy" on page B-2

- "wldeploy Ant Task Attribute Reference" on page B-4

Overview of the wldeploy Ant Task

The `wldeploy` Ant task enables you to perform `weblogic.Deployer` functions using attributes specified in an Ant XML file. You can use `wldeploy` along with other WebLogic Server Ant tasks to create a single Ant build script that:

- Builds your application from source, using `wlcompile`, `appc`, and the Web Services Ant tasks.

- Creates, starts, and configures a new WebLogic Server domain, using the `wlserver` and `wlconfig` Ant tasks.

- Deploys a compiled application to the newly-created domain, using the `wldeploy` Ant task.

See "Using Ant Tasks to Configure and Use a WebLogic Server Domain" on page 2-1 for more information about `wlserver` and `wlconfig`. See "Building Applications in a Split Development Directory" on page 4-1 for information about `wlcompile`.

Basic Steps for Using wldeploy

To use the `wldeploy` Ant task:

1. Set your environment.

   On Windows NT, execute the `setWLSEnv.cmd` command, located in the directory `WL_HOME\server\bin`, where `WL_HOME` is the top-level directory of your WebLogic Server installation.

   On UNIX, execute the `setWLSEnv.sh` command, located in the directory `WL_HOME/server/bin`, where `WL_HOME` is the top-level directory of your WebLogic Server installation.

2. In the staging directory, create the Ant build file (`build.xml` by default). If you want to use an Ant installation that is different from the one installed with WebLogic Server, start by defining the `wldeploy` Ant task definition:

   ```
   <taskdef name="wldeploy"
   classname="weblogic.ant.taskdefs.management.WLDeploy"/>
   ```

3. If necessary, add task definitions and calls to the `wlserver` and `wlconfig` tasks in the build script to create and start a new WebLogic Server domain. See "Using Ant Tasks to Configure and Use a WebLogic Server Domain" on page 2-1 for information about `wlserver` and `wlconfig`.

4. Add a call to `wldeploy` to deploy your application to one or more WebLogic Server instances or clusters. See "Sample build.xml Files for wldeploy" on page B-2 and "wldeploy Ant Task Attribute Reference" on page B-4.

5. Execute the Ant task or tasks specified in the `build.xml` file by typing `ant` in the staging directory, optionally passing the command a target argument:

   ```
   prompt> ant
   ```

Sample build.xml Files for wldeploy

The following example shows a `wldeploy` target that deploys an application to a single WebLogic Server instance:

```
<target name="deploy">
  <wldeploy
    action="deploy" verbose="true" debug="true"
    name="DeployExample" source="output/redeployEAR"
    user="weblogic" password="weblogic"
    adminurl="t3://localhost:7001" targets="myserver" />
</target>
```

The following example shows a corresponding task to undeploy the application; the example shows that when you undeploy or redeploy an application, you do not specify the source archive file or exploded directory, but rather, just its deployed name.:

```
<target name="undeploy">
  <wldeploy
    action="undeploy" verbose="true" debug="true"
    name="DeployExample"
    user="weblogic" password="weblogic"
    adminurl="t3://localhost:7001" targets="myserver"
    failonerror="false" />
</target>
```

The following example shows how to perform a partial redeploy of the application; in this case, just a single WAR file in the application is redeployed:

```
<target name="redeploy_partial">
  <wldeploy
    action="redeploy" verbose="true"
    name="DeployExample"
    user="weblogic" password="weblogic"
    adminurl="t3://localhost:7001" targets="myserver"
    deltaFiles="examples/general/redeploy/SimpleImpl.war" />
</target>
```

The following example uses the nested `<files>` child element of `wldeploy` to specify a particular file in the application that should be undeployed:

```
<target name="undeploy_partial">
  <wldeploy
    action="undeploy" verbose="true" debug="true"
    name="DeployExample"
    user="weblogic" password="weblogic"
```

```
        adminurl="t3://localhost:7001" targets="myserver"
        failonerror="false">
        <files
            dir="${current-dir}/output/redeployEAR/examples/general/redeploy"
            includes="SimpleImpl.jsp" />
    </wldeploy>
</target>
```

The following example shows how to deploy a Java EE library called `myLibrary` whose source files are located in the `output/myLibrary` directory:

```
<target name="deploy">
  <wldeploy action="deploy" name="myLibrary"
    source="output/myLibrary" library="true"
    user="weblogic" password="weblogic"
    verbose="true" adminurl="t3://localhost:7001"
    targets="myserver" />
</target>
```

wldeploy Ant Task Attribute Reference

The following sections describe the attributes and child element `<files>` of the `wldeploy` Ant task.

Main Attributes

The following table describes the main attributes of the `wldeploy` Ant task.

These attributes mirror some of the arguments of the `weblogic.Deployer` command. Oracle provides an Ant task version of the `weblogic.Deployer` command so that developers can easily deploy and test their applications as part of the iterative development process. Typically, however, administrators use the `weblogic.Deployer` command, and not the `wldeploy` Ant task, to deploy applications in a production environment. For that reason, see the weblogic.Deployer Command-Line Reference in *Deploying Applications to WebLogic Server* for

the full and complete definition of the attributes of the `wldeploy` Ant task. The table below is provided just as a quick summary.

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
action	The deployment action to perform. Valid values are `deploy`, `cancel`, `undeploy`, `redeploy`, `distribute`, `start`, and `stop`.	String
adminmode	Specifies that the deployment action puts the application into Administration mode. Administration mode restricts access to an application to a configured Administration channel. Valid values for this attribute are `true` and `false`. Default value is `false`, which means that by default the application is deployed in production mode so that all clients can access it immediately.	Boolean
adminurl	The URL of the Administration Server. The format of the value of this attribute is *protocol*`://`*host*`:`*port*, where *protocol* is either `http` or `t3`, *host* is the host on which the Administration Server is running, and *port* is the port which the Administration Server is listening. **Note:** In order to use the HTTP protocol, you must enable the http tunnelling option in the Administration Console.	String
allversions	Specifies that the action (redeploy, stop, and so on) applies to all versions of the application. Valid values for this attribute are `true` and `false`. The default value is `false`.	Boolean
altappdd	Specifies the name of an alternate Java EE deployment descriptor (`application.xml`) to use for deployment. If you do not specify this attribute, and you are deploying an Enterprise application, the default deployment descriptor is called `application.xml` and is located in the META-INF subdirectory of the main application directory or archive (specified by the `source` attribute.)	String

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
altwlsappdd	Specifies the name of an alternate WebLogic Server deployment descriptor (`weblogic-application.xml`) to use for deployment.	String
	If you do not specify this attribute, and you are deploying an Enterprise application, the default deployment descriptor is called `weblogic-application.xml` and is located in the META-INF subdirectory of the main application directory or archive (specified by the `source` attribute.)	
appversion	The version identifier of the deployed application.	String
debug	Enable `wldeploy` debugging messages.	Boolean
deleteFiles	Specifies whether to remove static files from a server's staging directory.	Boolean
	This attribute is valid only for unarchived deployments, and only for applications deployed using `stage` mode. You must specify target servers when using this attribute.	
	Specifying the `deleteFiles` attributes indicates that WebLogic Server should remove only thos files that it copied to the staging area during deployment.	
	This attribute can be used only in combination with `action="redeploy"`.	
	Because the `deleteFiles` attribute deletes all specified files, Oracle recommends that you use caution when using the `deleteFiles` attribute and that you do not use it in production environments.	
	Valid values for this attribute are true and false. Default value is false.	
deltaFiles	Specifies a comma- or space-separated list of files, relative to the root directory of the application, which are to be redeployed.	String
	Use this attribute only in conjunction with `action="redeploy"` to perform a partial redeploy of an application.	
enableSecurityValidation	Specifies whether or not to enable validation of security data.	Boolean
	Valid values for this attribute are true and false. Default value is false.	

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
externalStage	Specifies whether the deployment uses external_stage deployment mode. In this mode, the Ant task does not copy the deployment files to target servers; instead, you must ensure that deployment files have been copied to the correct subdirectory in the target servers' staging directories. You can specify only one of the following attributes: stage, nostage, or external_stage. If none is specified, the default deployment mode to Managed Servers is stage; the default mode to the Administration Server and in single-server cases is nostage. See Controlling Deployment File Copying with Staging Modes.	Boolean
failonerror	This is a global attribute used by WebLogic Server Ant tasks. It specifies whether the task should fail if it encounters an error during the build. Valid values for this attribute are true and false. Default value is true.	Boolean
graceful	Stops the application after existing HTTP clients have completed their work. You can use this attribute *only* when stopping or undeploying an application, or in other words, you must also specify either the action="stop" or action="undeploy" attributes. Valid values for this attribute are true and false. Default value is false.	Boolean
id	Identification used for obtaining status or cancelling the deployment. You assign a unique ID to an application when you deploy it, and then subsequently use the ID when redeploying, undeploying, stopping, and so on. If you do not specify this attribute, the Ant task assigns a unique ID to the application.	String
ignoresessions	This option immediately places the application into Administration mode without waiting for current HTTP sessions to complete. You can use this attribute *only* when stopping or undeploying an application, or in other words, you must also specify either the action="stop" or action="undeploy" attributes. Valid values for this attribute are true and false. Default value is false.	Boolean

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
libImplVer	Specifies the implementation version of a Java EE library or optional package.	String
	This attribute can be used only if the library or package does not include a implementation version in its manifest file. You can specify this attribute only in combination with the `library` attribute.	
	See "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.	
library	Identifies the deployment as a shared Java EE library or optional package. You must specify the `library` attribute when deploying or distributing any Java EE library or optional package.	Boolean
	Valid values for this attribute are `true` and `false`. Default value is `false`.	
	See "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.	
libSpecVer	Provides the specification version of a Java EE library or optional package.	String
	This attribute can be used only if the library or package does not include a specification version in its manifest file. You can specify this attribute only in combination with the `library` attribute.	
	See "Creating Shared Java EE Libraries and Optional Packages" on page 9-1.	
name	The deployment name for the deployed application.	String
	If you do not specify this attribute, WebLogic Server assigns a deployment name to the application, based on its archive file or exploded directory.	
nostage	Specifies whether the deployment uses nostage deployment mode.	Boolean
	In this mode, the Ant task does not copy the deployment files to target servers, but leaves them in a fixed location, specified by the `source` attribute. Target servers access the same copy of the deployment files.	
	You can specify only one of the following attributes: `stage`, `nostage`, or `external_stage`. If none is specified, the default deployment mode to Managed Servers is `stage`; the default mode to the Administration Server and in single-server cases is `nostage`.	
	See Controlling Deployment File Copying with Staging Modes.	

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
noversion	Indicates that the wldeploy Ant task should ignore all version related code paths on the Administration Server. This behavior is useful when deployment source files are located on Managed Servers (not the Administration Server) and you want to use the external_stage staging mode. If you use this option, you cannot use versioned applications. Valid values for this attribute are true and false. Default value is false.	Boolean
nowait	Specifies whether wldeploy returns immediately after making a deployment call (by deploying as a background task).	Boolean
password	The administrative password. **To avoid having the plain text password appear in the build file or in process utilities such as** ps, first store a valid username and encrypted password in a configuration file using the WebLogic Scripting Tool (WLST) storeUserConfig command. Then omit both the username and password attributes in your Ant build file. When the attributes are omitted, wldeploy attempts to login using values obtained from the default configuration file. If you want to obtain a username and password from a non-default configuration file and key file, use the userconfigfile and userkeyfile attributes with wldeploy. See the command reference for storeUserConfig in the WLST Command and Variable Reference for more information on storing and encrypting passwords.	String
plan	Specifies a deployment plan to use when deploying the application or module. By default, wldeploy does not use an available deployment plan, even if you are deploying from an application root directory that contains a plan.	String
planversion	The version identifier of the deployment plan.	String
remote	Specifies whether the server is located on a different machine. This affects how filenames are transmitted. Valid values for this attribute are true and false. Default value is false, which means that the Ant task assumes that all source paths are valid paths on the local machine.	Boolean

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
retiretimeout	Specifies the number of seconds before WebLogic Server undeploys the currently-running version of this application or module so that clients can start using the new version. It is assumed, when you specify this attribute, that you are starting, deploying, or redeploying a new version of an already-running application. See Updating Applications in a Production Environment.	int
securityModel	Specifies the security model to use for this deployment. Possible security models are: • Deployment descriptors only • Customize roles • Customize roles and policies • Security realm configuration (advanced model) Valid actual values for this attribute are DDOnly, CustomRoles, CustomRolesAndPolicy, or Advanced. See Options for Securing Web Application and EJB Resources for more information on these security models	String
source	The archive file or exploded directory to deploy.	File
stage	Specifies whether the deployment uses stage deployment mode. In this mode, the Ant task copies deployment files to target servers' staging directories. You can specify only one of the following attributes: stage, nostage, or external_stage. If none is specified, the default deployment mode to Managed Servers is stage; the default mode to the Administration Server and in single-server cases is nostage. See Controlling Deployment File Copying with Staging Modes.	Boolean
submoduletargets	Specifies JMS server targets for resources defined within a JMS application module. The value of this attribute is a comma-separated list of JMS server names. See the Using Sub-Module Targeting with JMS Application Modules.	String

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
targets	The list of target servers to which the application is deployed. The value of this attribute is a comma-separated list of the target servers, clusters, or virtual hosts. If you do not specify a target list when deploying an application, the target defaults to the Administration Server instance.	String
timeout	The maximum number of seconds to wait for a deployment to succeed.	int
upload	Specifies whether the source file(s) are copied to the Administration Server's upload directory prior to deployment. Use this attribute when you are on a remote machine and you cannot copy the deployment files to the Administration Server by other means. Valid values for this attribute are `true` and `false`. Default value is `false`.	Boolean
usenonexclusivelock	Specifies that the deployment action (deploy, redeploy, stop, and so on) uses the existing lock on the domain that has already been acquired by the same user performing the action. This attribute is particularly useful when the user is using multiple deployment tools (Ant task, command line, Administration console, and so on) simultaneously and one of the tools has already acquired a lock on the domain. Valid values for this attribute are `true` and `false`. Default value is `false`.	Boolean
user	The administrative username.	String
userconfigfile	Specifies the location of a user configuration file to use for obtaining the administrative username and password. Use this option, instead of the `user` and `password` attributes, in your build file when you do not want to have the plain text password shown in-line or in process-level utilities such as `ps`. Before specifying the `userconfigfile` attribute, you must first generate the file using using the WebLogic Scripting Tool (WLST) `storeUserConfig` command as described in the WLST Command and Variable Reference.	String

Table B-1 Attributes of the wldeploy Ant Task

Attribute	Description	Data Type
userkeyfile	Specifies the location of a user key file to use for encrypting and decrypting the username and password information stored in a user configuration file (the `userconfigfile` attribute). Before specifying the `userkeyfile` attribute, you must first generate the key file using the WebLogic Scripting Tool (WLST) `storeUserConfig` command as described in the WLST Command and Variable Reference.	String
verbose	Specifies whether `wldeploy` displays verbose output messages.	Boolean

Nested <files> Child Element

The `wldeploy` Ant task also includes the `<files>` child element that can be nested to specify a list of files on which to perform a deployment action (for example, a list of JSPs to undeploy.)

WARNING: Use of `<files>` to redeploy a list of files in an application has been deprecated as of release 9.0 of WebLogic Server. Instead, use the `deltaFiles` attribute of wldeploy.

The `<files>` element works the same as the standard `<fileset>` Ant task (except for the difference in actual task name). Therefore, see the Apache Ant Web site for detailed reference information about the attributes you can specify for the `<files>` element.

Spring Applications Reference

The following sections describe developing and managing Spring Framework-based applications for WebLogic Server. In some cases, the information in these sections is described from the perspective of creating MedRec-Spring.

Support for Spring on WebLogic Server

For information on how Oracle supports this release of WebLogic Server and the Spring Framework from SpringSource, see Supported Configurations for Products with Spring Framework.

Demonstrating Spring on WebLogic Server

To demonstrate the ways in which Spring can take advantage of WebLogic Server's enterprise features, Oracle redesigned the Avitek Medical Records sample application (MedRec) to replace core J2EE components with Spring components.

The following sections describe key steps that Oracle performed when redesigning MedRec. You can use this information if you want to redesign your own J2EE-based WebLogic Server applications to use Spring components. You can also leverage this information if you want to create a new application, based on Spring components, for WebLogic Server.

It is assumed that you are familiar with J2EE concepts, WebLogic Server, and the Spring Framework. For information on WebLogic Server, see WebLogic Server Documentation. For information on the Spring Framework in general, see http://www.springframework.org/.

Redesigning a J2EE-Based Application to a Spring-Based Application

To transform a J2EE-based application to a Spring-based application, you perform the following steps as desired:

1. Configure Spring Inversion of Control.

2. Enable the Spring Web Services Client Service. Spring offers a JAX-RPC factory which produces a proxy for Web Services.

3. Make JMS Services Available to the Application at Runtime.

4. Configure JMX: Expose the WebLogic Server Runtime MBean Server Connection to Spring.

5. Configure Spring JDBC to Communicate With the Connection Pool.

6. Use the Spring Transaction Abstraction Layer for Transaction Management.

7. Make Use of WebLogic Server Clustering and Clustered Spring Remoting.

The following sections describe the details of redesigning a J2EE-based application to a Spring-based application. Where appropriate, these sections include sample code. In most cases the sample code is from MedRec-Spring.

Configure Spring Inversion of Control

In Spring, references to other beans (injected properties) are configured via a Spring configuration XML file, `applicationContext-web.xml`.

In MedRec-Spring, Oracle replaced stateless session EJBs with POJOs in the Spring configuration file `src\medrecEar\web\WEB-INF\applicationContext-web.xml` as follows:

```
<bean name="/patient/record"
```

```
  class="com.bea.medrec.web.patient.actions.ViewRecordAction">
  <property name="medRecClientServiceFacade">
    <ref bean="medRecClientServiceFacade"/>
  </property>
</bean>
```

Then, in the application code, Oracle defined setter methods for the corresponding bean. For example:

```
protected MedRecClientServiceFacade medRecClientServiceFacade;

  public void setMedRecClientServiceFacade(

    MedRecClientServiceFacade pMedRecClientServiceFacade){

    this.medRecClientServiceFacade = pMedRecClientServiceFacade;

  }
```

Enable the Spring Web Services Client Service

To use Spring's JAX-RPC factory which produces a proxy for Web Services, you configure the Spring `JaxRpcPortProxyFactoryBean` by implementing code such as the following; in MedRec-Spring, Oracle implemented this code in the Spring configuration file `src\physicianEar\APP-INF\classes\applicationContext-phys-service.xml`.

```
<!-- reliable asynchronous web service for sending new medical records to
medrec -->
<bean id="reliableClientWebServicesPortType"
class="org.springframework.remoting.jaxrpc.JaxRpcPortProxyFactoryBean"
lazy-init="true">
<property name="wsdlDocumentUrl"
value="http://${WS_HOST}:${WS_PORT}/ws_phys/PhysicianWebServices?WSDL"/>
<property name="portName" value="PhysicianWebServicesPort"/>
<property name="jaxRpcService">
<ref bean="generatedReliableService"/>
</property>
<property name="serviceInterface"
value="com.bea.physician.webservices.client.PhysicianWebServicesPortType"/
>
<property name="username" value="medrec_webservice_user"/>
<property name="password" value="weblogic"/>
```

```
<property name="customProperties">
<props>
<prop key="weblogic.wsee.complex">true</prop>
</props>
</property>
</bean>

<> <!-- allows the jaxRpcService class to execute its constructor which
loads in type mappings -->
<bean id="generatedReliableService"
class="com.bea.physician.webservices.client.PhysicianWebServices_Impl">
</bean>
```

In this code example, note that:

- The `serviceInterface` represents Web Services operations.

- The `customProperties` property allows for custom WebLogic Server Web Service stub properties.

- The `jaxRpcService` value is set to WebLogic Server's generated JAX-RPC implementation service.

Make JMS Services Available to the Application at Runtime

In Spring, you must configure JMS services so that they are provided to the application during runtime. You can do this via a Spring bean that represents a messaging destination. In MedRec-Spring, Oracle made JMS services available to the application at runtime by implementing the following code in the Spring configuration file `src\medrecEar\APP-INF\classes\applicationContext-jms.xml`.

```
<bean id="uploadQueue"
class="org.springframework.jndi.JndiObjectFactoryBean">

    <property name="jndiName"

      value="com.bea.medrec.messagging.MedicalRecordUploadQueue"/>

  </bean>


  <bean id="jmsConnFactory"

    class="org.springframework.jndi.JndiObjectFactoryBean">
```

```
    <property name="jndiName"

      value="com.bea.medrec.messagging.MedRecQueueConnectionFactory"/>

  </bean>



  <bean id="uploadJmsTemplate"
class="org.springframework.jms.core.JmsTemplate">

    <property name="connectionFactory">

      <ref bean="jmsConnFactory"/>

    </property>

    <property name="defaultDestination">

      <ref bean="uploadQueue"/>

    </property>

  </bean>
```

Configure JMX: Expose the WebLogic Server Runtime MBean Server Connection to Spring

You can expose WebLogic Server's MBean Server to Spring through Spring's
`MBeanServerConnectionFactoryBean`, which is a convenience factory that produces an
`MBeanServerConnection` that is established and cached during application deployment and can
later be operated on by referencing beans. The `MBeanServerConnectionFactoryBean` can be
configured to return the WebLogic Server Runtime MBean Server, and to obtain a connection to
the WebLogic Server Domain Runtime MBean Server and the WebLogic Server Edit MBean
Server.

Note: Because the WebLogic Server Domain Runtime MBean Server is not active during
deployment, you must configure the `MBeanServerConnectionFactoryBean` to use
Spring's lazy instantiation. Lazy instantiation fetches the Spring bean when it is invoked.

Exposing the WebLogic Server Runtime MBean Server Connection to Spring is demonstrated in
the following code example, which, in MedRec-Spring, Oracle implemented in the Spring
configuration file `medrecEar/APP-INF/classes/applicationContext-jmx.xml`.

```
<> <!-- expose weblogic server's runtime mbeanserver connection -->
<bean id="runtimeMbeanServerConnection"
class="org.springframework.jmx.support.MBeanServerConnectionFactoryBean">
```

```
<property name="serviceUrl"
value="service:jmx:t3://${WS_HOST}:${WS_PORT}/jndi/weblogic.management.mbe
anservers.runtime"/>
<property name="environment">
<props>
<prop key="java.naming.security.principal">${WS_USERNAME}</prop>
<prop key="java.naming.security.credentials">${WS_USERNAME}</prop>
<prop
key="jmx.remote.protocol.provider.pkgs">weblogic.management.remote</prop>
</props>
</property>
</bean>
```

Configure Spring JDBC to Communicate With the Connection Pool

In MedRec-Spring, Oracle used a datasource that references a JDBC connection pool that is managed by WebLogic Server and also employed Spring's `JdbcDaoSupport` class. For information on `JdbcDaoSupport`, see the Spring documentation.

For an example of the way in which Oracle implemented JDBC, see the MedRec-Spring class

`src\medrecEar\dao\com\bea\medrec\dao\jdbc\JdbcPatientDao.java`

See also the following code examples, which, for MecRec-Spring, Oracle implemented in the Spring configuration files
`src\medrecEar\APP-INF\classes\applicationContext-db.xml` and
`src\medrecEar\APP-INF\classes\applicationContext-jdbc.xml`, respectively.

`applicationContext-db.xml` code example:

```
  <!-- datasource pool -->

  <bean id="dataSource"

    class="org.springframework.jndi.JndiObjectFactoryBean">

    <property name="jndiName" value="jdbc/MedRecGlobalDataSourceXA"/>

  </bean>
```

`applicationContext-jdbc.xml` code example:

```
  <bean id="patientDao"

    class="com.bea.medrec.dao.jdbc.JdbcPointBasePatientDao"
```

```
autowire="byType"/>
```

Additionally, in MedRec-Spring, Oracle replaced entity EJBs with POJOs and made use of Spring JDBC for persistence. For an example, see the MedRec-Spring class `\src\medrecEar\core\com\bea\medrec\domain\Address.java`

Use the Spring Transaction Abstraction Layer for Transaction Management

Spring supports distributed transactions through WebLogic Server's JTA implementation. You can also configure the Spring transaction manager to delegate responsibility to the WebLogic Server JTA transaction manager. This is accomplished via Spring's `WebLogicJtaTransactionManager` class. Oracle used this approach with MedRec-Spring in order to exactly mirror transaction management in the original version of MedRec.

To use the Spring transaction abstraction layer for transaction management and delegate responsibility to the WebLogic Server JTA transaction manager, you implement code such as the following, which Oracle implemented in the Spring configuration files `src\medrecEar\APP-INF\classes\applicationContext-tx.xml` and `src\medrecEar\APP-INF\classes\applicationContext-service.xml`, respectively.

`applicationContext-tx.xml` code example:

```
<!-- spring's transaction manager delegates to WebLogic Server's transaction
manager -->
<bean id="transactionManager"
class="org.springframework.transaction.jta.WebLogicJtaTransactionManager">
<property name="transactionManagerName"
value="javax.transaction.TransactionManager"/>
</bean>
```

`applicationContext-service.xml` code example:

```
<!-- base transaction proxy for which medrec spring beans inherit-->
< bean id="baseTransactionProxy"
class="org.springframework.transaction.interceptor.TransactionProxyFactory
Bean"
abstract="true">
<property name="transactionManager" ref="transactionManager"/>
<property name="transactionAttributes">
<props>
<prop key="activate*">PROPAGATION_REQUIRED</prop>
<prop key="create*">PROPAGATION_REQUIRED</prop>
```

```
<prop key="compose*">PROPAGATION_REQUIRED</prop>
<prop key="deny*">PROPAGATION_REQUIRED</prop>
<prop key="getRecord*">PROPAGATION_REQUIRED,readOnly</prop>
<prop key="getPatient*">PROPAGATION_REQUIRED,readOnly</prop>
<prop key="getLog*">PROPAGATION_NOT_SUPPORTED</prop>
<prop key="process*">PROPAGATION_REQUIRED</prop>
<prop key="save*">PROPAGATION_REQUIRED</prop>
<prop key="send*">PROPAGATION_REQUIRED</prop>
</props>
</property>
< /bean>

<!-- single point of service for all medrec clients -->
<bean id="medRecClientServiceFacade"
parent="baseTransactionProxy">
<property name="target">
<bean class="com.bea.medrec.service.MedRecClientServiceFacadeImpl">
<property name="adminService">
<ref bean="adminService"/>
</property>
<property name="patientService">
<ref bean="patientService"/>
</property>
<property name="recordService">
<ref bean="recordService"/>
</property>
<property name="recordXmlProcessorService">
<ref bean="recordXmlProcessorService"/>
</property>
</bean>
</property>
</bean>
```

The `transactionAttributes` you specify define the way in which Spring begins and ends transactions. Because MedRec-Spring delegates transaction management to WebLogic JTA, management tasks such as transaction suspension and rollback are handled as specified by WebLogic's transaction manager.

For more information on `WebLogicJtaTransactionManager`, see the Oracle Technology Network Web site.

Make Use of WebLogic Server Clustering

Spring applications can take advantage of WebLogic Server's clustering features. Because most Spring applications are packaged as Web applications (.war files), you need do not need to do anything special in order to take advantage of WebLogic Server clusters; all you need to do is deploy your Spring application to the servers in a WebLogic Server cluster.

Clustered Spring Remoting

The certification of Spring 1.2.8 and 2.0 on WebLogic Server extends the Spring `JndiRmiProxyFactoryBean` and its associated service exporter so that it supports proxying with any J2EE RMI implementation. To use the extension to the `JndiRmiProxyFactoryBean` and its exporter:

1. Configure client support by implementing code such as the following:

```
<bean id="proProxy"
class="org.springframework.remoting.rmi.JndiRmiProxyFactoryBean">
<property name="jndiName" value="t3://${serverName}:${rmiPort}/order"/>
</property>
<property name="jndiEnvironment">
<props>
<prop key="java.naming.factory.url.pkgs">weblogic.jndi.factories</prop>
</props>
</property>
<property name="serviceInterface"
value="org.springframework.samples.jpetstore.domain.logic.OrderService"
/>
</bean>
```

2. Configure the service exporter by implementing code such as the following:

```
<bean id="order-pro"
class="org.springframework.remoting.rmi.JndiRmiServiceExporter">
<property name="service" ref="petStore"/>
<property name="serviceInterface"
value="org.springframework.samples.jpetstore.domain.logic.OrderService"
/>
<property name="jndiName" value="order"/>
</bean>
```

WebLogic Spring Security Integration

The WebLogic Server security system supports and extends Java EE security while providing a rich set of security providers that you can be customize to integrate with different security databases or security policies.

As described at the Spring Security Web site, Acegi Security is now Spring Security, the official security project of the Spring Portfolio. The Spring security (acegi) framework provides security to a Spring application and includes a rich set of security providers.

The question then becomes how to integrate the two security frameworks.

For a combined J2EE and Spring application, rather than require authentication with both security frameworks, WLS security and Spring security work together. WLS security handles the authentication via the default Authentication provider for the security realm, and converts WLS principals to Spring GrantedAuthority principals through a mapper class. Once authenticated by WLS security, a user is authenticated for Spring security. You can then decide how to secure the objects in the application. One common practice is to secure Java EE resource with Weblogic security and secure Spring resource with Spring security.

Spring Container Adapter Provides Integration

As described in the Spring Security Reference, Container Adapters enable Acegi Security to integrate directly with the containers used to host end user applications, in this case WebLogic Server.

The integration between a container and Acegi Security is achieved through an adapter. The adapter provides a container-compatible user authentication provider, and needs to return a container-compatible user object.

applicationContext-acegi-security.xml is the configuration file for Spring security. For WebLogic Server, WeblogicAuthenticationFilter is added to the list of filters in applicationContext-acegi-security.xml. This filter is responsible for converting the Weblogic principals to Spring GrantedAuthority subjects, based on the mapper. The mapper is configured as a property for the WeblogicAuthenticationFilter, and it is injected at creation time.

The following is an example of the mapper class.

```
public class MyAuthorityGranter implements AuthorityGranter {
public Set grant(Principal principal) {
Set rtnSet = new HashSet();
if (principal.getName().equals("fred@oracle.com")) {
```

```
rtnSet.add("ROLE_SUPERVISOR");
rtnSet.add("IS_AUTHENTICATED_ANONYMOUSLY");
}
return rtnSet;
}
}
```

In this example, user `fred@oracle.com` in the WebLogic domain is mapped to `ROLE_SUPERVISOR` and `IS_AUTHENTICATED_ANONYMOUSLY`.

How applicationContext-acegi-security.xml Is Plugged Into web.xml

The following code is added to `web.xml` to plug in the `applicationContext-acegi-security.xml` file:

```
<filter>
<filter-name>Acegi Filter Chain Proxy</filter-name>
<filter-class>org.acegisecurity.util.FilterToBeanProxy</filter-class>
<init-param>
<param-name>targetClass</param-name>
<param-value>org.acegisecurity.util.FilterChainProxy</param-value>
</init-param>
</filter>

<filter-mapping>
<filter-name>Acegi Filter Chain Proxy</filter-name>
<url-pattern>/main/secure/*</url-pattern>
</filter-mapping>

<listener>
<listener-class>org.springframework.web.context.ContextLoaderListener</lis
tener-
class>
</listener>

<context-param>
<param-name>contextConfigLocation</param-name>
<param-value>
/WEB-INF/applicationContext-acegi-security.xml
```

```
        </param-value>
    </context-param>
```

Spring Extension to WebLogic Server

A standard WebLogic Server installation provides standard Java EE 5 dependency injection and interceptors (a form of aspect-oriented programming) in the WebLogic Server Java EE container. WebLogic Server also supports a Spring Framework extension that provides enhanced dependency injection and aspect-oriented programming features in the container. This extension uses Pitchfork, a Spring Framework add-on that provides JSR-250 (Common Annotations), dependency injection, and EJB 3.0 style interception. The extension provides dependency injection and aspect-oriented programming to EJB instances and Web components that include the servlet listener and filter.

Note: JSP tag handlers do not support the Spring extension in this release of WebLogic Server.

To enable the Spring extension with WebLogic Server, do the following:

1. Download a version of Spring and its dependencies. Download the version of Spring that is certified by Oracle.

    You must have at least the following jar files:

    – spring.jar

    – aspectjweaver.jar

    – commons-logging.jar

    – log4j-1.2.14.jar

    –   pitchfork.jar

    You can add other jar files if necessary.

2. Add the jars listed above to the WebLogic Server classpath.

    The WebLogic Server Web container and EJB container use these jars to provide container service (dependency injection and interceptor).

    Applications (packaged as ear, war, or jar files) use these jars because they are in the server classpath. You can configure your application to use the version of jars packaged with the application if you enable certain of the deployment descriptors.

3. Enable the Spring extension by setting the `<component-factory-class-name>` element to `org.springframework.jee.interfaces.SpringComponentFactory`. This element

exists in EJB, Web, and application descriptors. A module level descriptor overwrites an application level descriptor. If the tag is set to null (default), the Spring extension is disabled.

4. Provide the common Spring bean definition file with the name `spring-ejb-jar.xml` or `spring-web.xml`, and place it in the /WEB-INF/classes/META-INF directory of your application (or put the META-INF directory in a jar file). These are the standard Spring bean definition files with the names that the Weblogic container searches for. For the Spring container to be aware of the EJB or servlet instance, the `<id>` tag of the Spring bean must be set to the `ejb-name` for EJB or the `class-name` of the web components.

Spring Console Extension

You can use an extension to the WebLogic Server Administration Console to monitor Spring bean attributes that are defined in applications.

The Spring console extension is based on RuntimeMBeans registered using the WebLogic Server infrastructure. The console extension displays configuration information for deployed Spring beans.

In the current release, the Spring console extension works with Web applications but not EJBs.

To use the Spring console extension, you must turn on support for Spring beans and enable the Spring console extension, as follows:

1. Deploy *WL_HOME*/server/lib/weblogic-spring.jar to WebLogic Server, where *WL_HOME* refers to the main WebLogic Server installation directory, such as \beahome/wlserver_10.3. You only need to perform this step once for your WebLogic Server instance.

   This jar file is a Java EE optional package used by an application (packaged as an ear or war file) to create the MBeans for the application during its deployment.

   Deploy `weblogic-spring.jar` either of the following ways:

   – Use the Administration Console, as described in Deploy applications and modules in the *Administration Console Online Help*. The jar file is located in *WL_HOME*\server\lib.

   – Issue the following command at the command line:

   ```
   java weblogic.Deployer -library -deploy –source
      WL_HOME/server/lib/weblogic-spring.jar –targets
      server_name -adminurl server_URL -user
      user_name -password password
   ```

2. Enable the Spring console extension in the Administration Console. (The Spring console extension is disabled by default.) You only need to perform this step once for your domain. Do the following:

   a. Log into the Administration Console.

   b. In the banner toolbar region at the top of the right pane of the Console, click **Preferences**.

   c. On the Preferences page, click **Extensions**.

   d. Select the check box next to **spring-console**, then click **Enable**.

   e. Stop the server, then restart it for the change to take effect.

3. Change the manifest of your application (packaged as an ear or war file) so it includes `weblogic-spring.jar` as a Java EE optional package. Do this to each Spring application you want to make use of the Spring runtime MBeans or Spring console extension

   Do this by adding the following lines to your `META-INF/Manifest.mf`:

   ```
   Extension-List: WeblogicSpring

   WeblogicSpring-Extension-Name: weblogic-spring

   WeblogicSpring-Specification-Version: 10.3.0.0

   WeblogicSpring-Implementation-Version: 10.3.0.0
   ```

4. To generate runtime MBeans for your application, the WeblogicSpringApplicationListener from `weblogic-spring.jar` must be added as a listener to the root Spring context. This can be done in two ways. The first way is preferred.

   a. Edit `web.xml` to change the Spring context loader listener to the WebLogic context loader listener. That is, change this:

   ```
   <listener-class>
       org.springframework.web.context.ContextLoaderListener
   </listener-class>
   ```

   to this:

   ```
   <listener-class>
       weblogic.spring.monitoring.WeblogicContextLoaderListener
   </listener-class>
   ```

   WeblogicContextLoaderListener extends the Spring ContextLoaderListener and adds WeblogicSpringApplicationListener as an ApplicationListener and a BeanFactoryPostProcessor to the Web application context.

b. If you already have a customized ContextLoaderListener and cannot use the Weblogic one, you can add WeblogicSpringApplicationListener as a bean to the root ApplicationContext.

In your Web application Spring XML configuration file add:

```
<bean
class="weblogic.spring.monitoring.WeblogicSpringApplicationListener
" />
```

Exposing Spring Beans Through the WebLogic Administration Console

To access Spring beans that are not MBeans through the WebLogic Administration Console, you must configure an MBeanExporter in the `applicationContext.xml` file and specify which beans to expose via the assembler. Make sure that the `applicationName` property is the deployed name of your application.