

Oracle® WebLogic Server

Deploying Applications to WebLogic Server

10g Release 3 (10.3)

July 2008

ORACLE®

Copyright © 2007, 2008, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Introduction and Roadmap

Document Scope and Audience	1-1
Guide to This Document	1-2
Standards Compatibility	1-3
Deployment Terminology	1-3
Related Documentation	1-4
New and Changed Features in This Release	1-5

Understanding WebLogic Server Deployment

Overview of the Deployment Process	2-1
J2EE 1.4 Deployment Implementation	2-2
WebLogic Server Deployment Features	2-4
Additional Deployment Configuration Properties	2-4
Exporting Applications for Deployment to Multiple Environments	2-4
Administration Mode for Isolating Production Applications	2-5
Deployable JDBC, JMS, and WLDF Application Modules	2-5
Module-Level Deployment and Redeployment for Enterprise Applications	2-5
Safe Redeployment for Production Applications	2-6
Security Roles Required for Deployment	2-6
Supported Deployment Units	2-6
Enterprise Application	2-7
Web Application	2-7

Enterprise JavaBean	2-8
Resource Adapter	2-8
Web Service	2-8
J2EE Library	2-8
Optional Package	2-9
JDBC, JMS, and WLDF Modules	2-9
Client Application Archive	2-10
Deployment Tools	2-10
weblogic.Deployer	2-11
Administration Console	2-11
WLST	2-11
Deployment Tools for Developers	2-11

Preparing Applications and Modules for Deployment

Packaging Files for Deployment	3-2
Using Archived Files	3-2
Using Exploded Archive Directories	3-3
Java EE Rules for Deploying Exploded EAR Directories without Deployment Descriptors	3-3
Creating an Exploded Archive Directory from an Archive File	3-4
Understanding Default Deployment Names	3-4
Understanding Application Naming Requirements	3-5
Understanding Deployment Version Strings	3-6
Creating an Application Installation Directory	3-6
Steps for Creating an Application Installation Directory	3-7
Using FastSwap Deployment to Minimize Redeployment	3-10
How FastSwap Deployment Works	3-10
Supported FastSwap Application Configurations	3-10

Enabling FastSwap In Your Application	3-11
Overview of the FastSwap Process	3-11
Using Ant with the JMX Interface	3-12
Application Types and Changes Supported with FastSwap	3-13
Limitations When Using FastSwap	3-17
Handling Unsupported FastSwap Changes	3-17
Best Practices for Preparing Deployment Files	3-19

Configuring Applications for Production Deployment

Understanding the Deployment Configuration Process	4-1
Deployment Configuration Lifecycle.	4-2
Understanding Application Deployment Descriptors	4-3
Understanding WebLogic Server Deployment Plans	4-4
Goals for Production Deployment Configuration	4-5
Typical Deployment Configuration Workflows	4-6
Application with Single Deployment Plan	4-6
Single Deployment Plan Ownership and Limitations	4-9
Application with Multiple Deployment Plans	4-9
Benefits of a Multiple Deployment Plan Workflow	4-11
Multiple Deployment Plan Ownership and Limitations	4-11
Creating a New Deployment Plan to Configure an Application	4-11
Preparing the Deployment Files	4-12
Installing the Application Archive	4-12
Saving Configuration Changes to a Deployment Plan	4-13
Understanding Deployment Plan Contents	4-13
Using an Existing Deployment Plan to Configure an Application	4-17
Generic File Loading Overrides	4-19
How It Works	4-19

Directory Structure	4-20
Application Usage	4-20
Additional Configuration Tasks	4-21
Best Practices for Managing Application Configuration	4-21

Exporting an Application for Deployment to New Environments

Overview of the Export Process	5-1
Goals for Exporting a Deployment Configuration.	5-2
Tools for Exporting a Deployment Configuration	5-2
Understanding Deployment Property Classifications	5-3
Steps for Exporting an Application's Deployment Configuration	5-4
Staging Application Files for Export	5-4
Generating a Template Deployment Plan using weblogic.PlanGenerator	5-5
Customizing the Deployment Plan Using the Administration Console	5-6
Install the Exported Application and Template Deployment Plan.	5-7
Add Variables for Selected Tuning Properties.	5-7
Retrieve the Customized Deployment Plan	5-7
Manually Customizing the Deployment Plan	5-7
Removing Variables from a Deployment Plan.	5-8
Assigning Null Variables to Require Administrator Input.	5-8
Validating the Exported Deployment Configuration	5-8
Best Practices for Exporting a Deployment Configuration.	5-9

Deploying Applications and Modules with weblogic.deployer

Overview of Common Deployment Scenarios	6-2
Uploading Deployment Files from a Remote Client.	6-3
Upload Behavior When Updating an Application or Plan.	6-3
Deploying to a Single-Server Domain	6-4

Deploying an Application with a Deployment Plan	6-4
Deploying an Application That Looks Up System Resources from JNDI During preStart6-5	
Targeting Deployments to Servers, Clusters, and Virtual Hosts	6-6
Understanding Deployment Targets	6-6
Deploying to One or More Targets	6-7
Deploying to a Cluster Target	6-8
Enforcing Consistent Deployment to All Configured Cluster Members	6-8
Using Module-Level Targeting for Deploying an Enterprise Application	6-9
Module-Targeting Syntax	6-10
Targeting Web Application Modules	6-10
Deploying JDBC, JMS, and WLDF Application Modules	6-11
Targeting Application-Scoped JMS, JDBC, and WLDF Modules	6-11
Using Sub-Module Targeting with JMS Application Modules	6-12
Default Submodule Targeting	6-12
Sub-module Targeting for Stand-alone JMS Modules	6-13
Sub-module Targeting for Application-scoped JMS Modules	6-13
Controlling Deployment File Copying with Staging Modes	6-14
Staging Mode Descriptions and Best Practices	6-15
Using Nostage Mode Deployment	6-17
Syntax for Nostage Mode	6-18
Using Stage Mode Deployment	6-18
Syntax for Stage Mode	6-19
Using External_stage Mode Deployment	6-19
Syntax for external_stage Mode	6-20
Changing the Default Staging Behavior for a Server	6-21
Distributing Applications to a Production Environment	6-21
Distributing an Application	6-22
Starting a Distributed Application in Administration Mode	6-22

Starting a Distributed Application	6-22
Deploying Shared Java EE Libraries and Dependent Applications	6-22
Understanding Deployment Behavior for Shared Libraries	6-23
Registering Libraries with WebLogic Server.	6-24
Specifying Library Versions at Deployment	6-24
Deploying Applications That Reference Libraries	6-25
Best Practices for Deploying Applications	6-26

Auto-Deploying Applications in Development Domains

Enabling and Disabling Auto-Deployment.	7-2
Auto-Deploying, Redeploying, and Undeploying Archived Applications	7-2
Auto-Deploying, Redeploying, and Undeploying Exploded Archives.	7-3

Redeploying Applications in a Production Environment

Overview of Redeployment Strategies	8-1
Production Redeployment	8-2
In-Place Redeployment	8-2
Partial Redeployment of Static Files.	8-3
Partial Redeployment of J2EE Modules.	8-3
Understanding When to Use Different Redeployment Strategies.	8-4
Using Production Redeployment to Update Applications	8-5
How Production Redeployment Works	8-5
Production Redeployment In Clusters	8-6
Requirements and Restrictions for Production Redeployment	8-6
Development Requirements	8-6
Deployment Requirements	8-8
Restrictions on Production Redeployment Updates	8-8
Specifying an Application Version Identifier.	8-8

Assigning a Version Identifier During Deployment and Redeployment	8-8
Displaying Version Information for Deployed Applications	8-9
Redeploying a New Version of an Application	8-9
Undeploying a Retiring Application	8-11
Rolling Back the Production Redeployment Process	8-11
Graceful Shut Down of RMI Client Request Processing	8-12
Distributing a New Version of a Production Application	8-13
Steps for Distributing a New Version of an Application	8-14
Making an Application Available to Clients	8-14
Best Practices for Using Production Redeployment	8-15
Using In-Place Redeployment for Applications and Stand-alone Modules	8-16
Redeploying Applications and Modules In-Place	8-18
Best Practices for Redeploying Applications and Modules In-Place	8-19
Using Partial Redeployment for J2EE Module Updates	8-19
Restrictions for Updating J2EE Modules in an EAR	8-20
Best Practices for Updating J2EE Modules in an EAR	8-20
Updating Static Files in a Deployed Application	8-21
Updating the Deployment Configuration for an Application	8-22
Modifying a Configuration Using the Administration Console	8-22
How Configuration Changes Are Stored	8-22
Updating an Application to Use a Different Deployment Plan	8-22
Understanding Redeployment Behavior for Deployment Configuration Changes	8-23

Managing Deployed Applications

Taking a Production Application Offline	9-1
Stopping an Application to Restrict Client Access	9-2
Undeploying an Application or Module	9-2
Undeploying Shared Libraries and Packages	9-3

Adding a New Module to a Deployed Enterprise Application	9-3
Changing the Order of Deployment at Server Startup	9-4
Changing the Deployment Order for Applications and Stand-alone Modules	9-4
Changing the Deployment Order for Modules in an Enterprise Application	9-5
Ordering Startup Class Execution and Deployment	9-5
Changing the Target List for an Existing Deployment	9-6
Removing Files from a Web Application Deployment	9-6
Managing Long-Running Deployment Tasks	9-7
On-demand Deployment of Internal Applications	9-8
Internal Application Types	9-8
Configuring On-Demand Deployment	9-8
Configuring On-Demand Deployment Using the Administration Console	9-8
Configuring On-Demand Deployment Using WLST	9-9

weblogic.Deployer Command-Line Reference

Required Environment for weblogic.Deployer	A-1
Syntax for Invoking weblogic.Deployer	A-2
Connection Arguments	A-4
User Credentials Arguments	A-4
Common Arguments	A-5
Commands and Options	A-7
Cancel	A-8
Deploy	A-9
Distribute	A-13
Listapps	A-16
List, Listtask	A-17
Redeploy	A-17
Start	A-21

Stop	A-23
Undeploy	A-25
Update	A-27
Example config.xml File and Corresponding weblogic.Deployer Command	A-29

weblogic.PlanGenerator Command Line Reference

Overview of weblogic.PlanGenerator	B-1
Required Environment for weblogic.PlanGenerator	B-2
Syntax for Invoking weblogic.PlanGenerator	B-2
Options	B-3
Common weblogic.PlanGenerator Tasks	B-4
Creating an Initial Deployment Plan in an Application's Root Directory	B-4
Creating a New Deployment Plan Based on an Existing Plan	B-5
Controlling the Components Exported to a Deployment Plan	B-5

Introduction and Roadmap

The following sections describe the contents and organization of this guide—*Deploying Applications to WebLogic Server*:

- [“Document Scope and Audience”](#) on page 1-1
- [“Guide to This Document”](#) on page 1-2
- [“Standards Compatibility”](#) on page 1-3
- [“Deployment Terminology”](#) on page 1-3
- [“Related Documentation”](#) on page 1-4
- [“New and Changed Features in This Release”](#) on page 1-5

Document Scope and Audience

This document is a resource for:

- Administrators who want to deploy Java 2 Platform, Enterprise Edition (J2EE) applications or application modules to WebLogic Server® instances or clusters. It is assumed that you are working in a production environment, which is generally characterized by multiple WebLogic Server instances or clusters running on multiple machines. It is also assumed that you have one or more application module archive files that have been tested and are ready to deploy on a production server.
- Developers who may need to deploy an application in a development environment, package an application for delivery to an administrator or deployer, or create and export

the configuration of an application for deployment to a testing, staging, or production environment

Guide to This Document

- This chapter, [Introduction and Roadmap](#), describes the organization of this document and highlights new deployment features introduced in this release of WebLogic Server.
- [Chapter 2, “Understanding WebLogic Server Deployment,”](#) provides an overview of deployment features used in WebLogic Server.
- [Chapter 3, “Preparing Applications and Modules for Deployment,”](#) explains how to prepare application and module files for deployment to WebLogic Server.
- [Chapter 4, “Configuring Applications for Production Deployment,”](#) how to configure an application for deployment to a specific WebLogic Server environment.
- [Chapter 5, “Exporting an Application for Deployment to New Environments,”](#) describes how to export an application's WebLogic Server deployment configuration to a custom deployment plan, which administrators use in deploying the application into non-development environments.
- [Chapter 6, “Deploying Applications and Modules with `weblogic.deployer`,”](#) describes basic and advanced techniques for deploying applications to Weblogic Server.
- [Chapter 7, “Auto-Deploying Applications in Development Domains,”](#) describes how to quickly deploy an application to a stand-alone server (Administration Server) for evaluation or testing in a single-server development environment.
- [Chapter 8, “Redeploying Applications in a Production Environment,”](#) explains how to safely update, redeploy, and reconfigure applications that you have deployed to a production environment.
- [Chapter 9, “Managing Deployed Applications,”](#) describes common tasks that an Administrator performs when managing deployed applications and modules.
- [Appendix A, “`weblogic.Deployer` Command-Line Reference,”](#) provides a complete reference for the `weblogic.Deployer` tool syntax.
- [Appendix B, “`weblogic.PlanGenerator` Command Line Reference,”](#) describes how to use the `weblogic.PlanGenerator` utility to create a basic deployment plan.

Standards Compatibility

WebLogic Server implements the J2EE 1.4 specification. J2EE 1.4 includes a deployment specification, JSR-88, that describes a standard API used by deployment tools and application server providers to configure and deploy applications to an application server.

WebLogic Server implements both the JSR-88 Service Provider Interface (SPI) plug-in and model plug-in to comply with the J2EE 1.4 deployment specification. You can use a basic J2EE 1.4 deployment API deployment tool with the WebLogic Server plug-ins (without using WebLogic Server extensions to the API) to configure, deploy, and redeploy J2EE applications and modules to WebLogic Server. See [Compatibility Statement for WebLogic Server](#).

Deployment Terminology

The following WebLogic Server deployment terms are used throughout this document:

- *application*—One or more software programs, used collectively by an end user to perform computing tasks.
- *application installation directory*—A WebLogic Server directory structure designed to help organize deployment files and generated deployment configuration artifacts for an application or module. Also referred to as an *application root directory*.
- *application module*—An XML document used to configure JMS or JDBC resources. An application module can be one of the following types:
 - *stand-alone*—Resources are bound to the global JNDI tree.
 - *application-scoped*—Bundled as part of an Enterprise Application and scoped within the application itself.
- *application version*—A string value that identifies the version of a deployed application. Compatible applications that use version strings can use the WebLogic Server production redeployment strategy.
- *deployment configuration*—The process of defining the deployment descriptor values required to deploy an application to a particular WebLogic Server domain. The deployment configuration for an application or module is stored in three types of XML document: J2EE deployment descriptors, WebLogic Server descriptors, and WebLogic Server deployment plans.
- *deployment descriptor*—An XML document used to define the J2EE behavior or WebLogic Server configuration of an application or module at deployment time.

- *deployment plan*—An XML document used to define an application’s WebLogic Server deployment configuration for a specific WebLogic Server environment, such as development, test, or production. A deployment plan resides outside of an application’s archive file and contains deployment properties which override an application’s existing WebLogic Server deployment descriptors. Use deployment plans to easily change an application’s WebLogic Server configuration for a specific environment *without* modifying existing deployment descriptors. Multiple deployment plans can be used to reconfigure a single application for deployment to multiple, differing WebLogic Server environments.
- *distribution*—The process by which WebLogic Server copies deployment source files to target servers for deployment.
- *production redeployment*—A WebLogic Server redeployment strategy that deploys a new version of a production application alongside an older version, while automatically managing HTTP connections to ensure uninterrupted client access.
- *staging mode*—The method WebLogic Server uses to make deployment files available to target servers in a domain. Staging modes determine whether or not files are distributed (copied) to target servers before deployment.

Related Documentation

For additional information about deploying applications and modules to WebLogic Server, see these documents:

- [Programming WebLogic Deployment](#) describes the WebLogic Server deployment API, which implements and extends the J2EE 1.4 specification. All WebLogic Server deployment tools use this API.
- [Developing WebLogic Server Applications](#) describes how to deploy applications during development using the `wldeploy` Ant task, and provides information about the WebLogic Server deployment descriptor for Enterprise Applications.
- The WebLogic Server J2EE programming guides describe the J2EE and WebLogic Server deployment descriptors used with each J2EE application and module:
 - [Developing Web Applications, Servlets, and JSPs for WebLogic Server](#)
 - [Programming WebLogic Enterprise JavaBeans \(EJB\)](#)
 - [Programming WebLogic Server Resource Adapters](#)
 - [Programming WebLogic Web Services for WebLogic Server](#)
 - [Programming WebLogic Deployment](#)

- *Programming WebLogic JDBC* describes the XML deployment descriptors for JDBC application modules.
- *Programming WebLogic JMS* describes the XML deployment descriptors for JMS application modules.

New and Changed Features in This Release

For a comprehensive listing of the new WebLogic Server features introduced in this release, see [“What’s New in WebLogic Server”](#) in *Release Notes*.

Introduction and Roadmap

Understanding WebLogic Server Deployment

The following sections provide an overview of WebLogic Server deployment:

- [“Overview of the Deployment Process” on page 2-1](#)
- [“J2EE 1.4 Deployment Implementation” on page 2-2](#)
- [“WebLogic Server Deployment Features” on page 2-4](#)
- [“Supported Deployment Units” on page 2-6](#)
- [“Deployment Tools” on page 2-10](#)

Overview of the Deployment Process

The term *application deployment* refers to the process of making an application or module available for processing client requests in a WebLogic Server domain. Application deployment generally involves the following tasks:


- [“Preparing Applications and Modules for Deployment” on page 3-1](#)
- [“Configuring Applications for Production Deployment” on page 4-1](#)
- [“Exporting an Application for Deployment to New Environments” on page 5-1](#)
- [“Deploying Applications and Modules with `weblogic.deployer`” on page 6-1](#)
- [“Redeploying Applications in a Production Environment” on page 8-1](#)
- [“Managing Deployed Applications” on page 9-1](#)

J2EE 1.4 Deployment Implementation

WebLogic Server implements the J2EE 1.4 specification. J2EE 1.4 includes a deployment specification, JSR-88, that describes a standard API used by deployment tools and application server providers to configure and deploy applications to an application server.

WebLogic Server implements both the JSR-88 Service Provider Interface (SPI) plug-in and model plug-in to comply with the J2EE 1.4 deployment specification. You can use a basic J2EE 1.4 deployment API deployment tool with the WebLogic Server plug-ins (without using WebLogic Server extensions to the API) to configure, deploy, and redeploy J2EE applications and modules to WebLogic Server. The WebLogic Server configuration generated by a J2EE 1.4 deployment API configuration process is stored in a deployment plan and one or more generated WebLogic Server deployment descriptor files, as shown in [Figure 2-1](#). WebLogic Server deployment descriptors are generated as needed to store WebLogic Server configuration data.

Figure 2-1 Configuring Applications with the J2EE 1.4 Deployment API



The WebLogic Server deployment plan generated by a J2EE 1.4 deployment API deployment tool identifies the WebLogic Server deployment descriptors that were generated for the application during the configuration session.

Although the J2EE 1.4 deployment API provides a simple, standardized way to configure applications and modules for use with a J2EE 1.4-compliant application server, the specification does not address many deployment features that were available in previous WebLogic Server releases. For this reason, WebLogic Server provides important extensions to the J2EE 1.4 deployment API specification to support capabilities described in [“WebLogic Server Deployment Features”](#) on page 2-4.

WebLogic Server Deployment Features

WebLogic Server supports the following advanced deployment features to help you reliably deploy and manage applications in a production environment.

- [“Additional Deployment Configuration Properties” on page 2-4](#)
- [“Exporting Applications for Deployment to Multiple Environments” on page 2-4](#)
- [“Administration Mode for Isolating Production Applications” on page 2-5](#)
- [“Deployable JDBC, JMS, and WLDF Application Modules” on page 2-5](#)
- [“Module-Level Deployment and Redeployment for Enterprise Applications” on page 2-5](#)
- [“Safe Redeployment for Production Applications” on page 2-6](#)
- [“Security Roles Required for Deployment” on page 2-6](#)

Additional Deployment Configuration Properties

Whereas the J2EE 1.4 deployment API deployment specification enables you to generate vendor-specific descriptor values necessary for deploying an application, WebLogic Server extensions to J2EE 1.4 deployment API allow you to configure many additional deployment properties, including:

- The names of external resources required for the application to operate
- The declared names of services provided in a deployed application (JNDI names), which other applications may reference for their own use
- Tuning properties that control the performance and behavior of the application on WebLogic Server

You can store these deployment properties in WebLogic Server deployment plans.

Exporting Applications for Deployment to Multiple Environments

The basic J2EE 1.4 deployment API configuration process provides a simple way for standardized deployment tools to deploy J2EE applications on multiple application server products. However, it does not help in the process of migrating an application’s configuration from one environment to another within an organization. The WebLogic Server deployment API

extends the J2EE 1.4 deployment API to provide support for exporting an application's configuration for deployment to multiple WebLogic Server environments, such as testing, staging, and production domains. See [“Exporting an Application for Deployment to New Environments” on page 5-1](#).

Administration Mode for Isolating Production Applications

Distributing an application copies deployment files to target servers and places the application in a prepared state. You can then start the application in Administration mode, which restricts access to the application to a configured Administration channel so you can perform final testing without opening the application to external client connections or disrupting connected clients. You can start an application in administration mode with the `-adminmode` option as described in [“Starting a Distributed Application in Administration Mode” on page 6-22](#). See [“Distributing Applications to a Production Environment” on page 6-21](#) and [“Distributing a New Version of a Production Application” on page 8-13](#).

After performing final testing, you can either undeploy the application to make further changes, or start the application in Production mode to make it generally available to clients.

See [“Distributing an Application” on page 6-22](#).

Deployable JDBC, JMS, and WLDF Application Modules

JDBC, JMS, and WLDF resources can be stored as application modules, which can be deployed stand-alone to multiple servers or clusters or included within an Enterprise Application as application-scoped resources. Stand-alone JDBC, JMS, and WLDF application modules make it easy to replicate resources in multiple WebLogic Server domains. Application-scoped resource modules make it possible to include all of an application's required resources within the application module itself, for maximum portability to multiple environments. See [Developing Applications with WebLogic Server](#) for more information about using application-scoped resources. See [“Deploying JDBC, JMS, and WLDF Application Modules” on page 6-11](#) to deploy stand-alone or application-scoped resources to WebLogic Server.

Module-Level Deployment and Redeployment for Enterprise Applications

WebLogic Server enables you to target individual modules of an Enterprise Application to different server targets, or to deploy only a subset of available modules in an Enterprise Application. This provides flexible packaging options, allowing you to bundle a group of related

modules together in an Enterprise Application, but deploy only selected modules to individual servers in a domain.

Safe Redeployment for Production Applications

WebLogic Server enables you to safely update and redeploy a new version of a production application without affecting current HTTP clients to the application. Production redeployment helps you roll out bug fixes or new functionality without application downtime, and without creating redundant servers in order to roll out the changes. See [“Redeploying Applications in a Production Environment” on page 8-1](#).

Security Roles Required for Deployment

The built-in security roles for “Admin” and “Deployer” users allow you to perform deployment tasks using the WebLogic Server Administration Console. The “AppTester” security role allows you to test versions of applications that are deployed to Administration mode. When deploying across WebLogic domains, the “CrossDomainConnector” role allows you to make inter-domain calls from foreign domains. For a complete listing of all security roles, see [Default Global Roles](#) in *Securing WebLogic Resources Using Roles and Policies*.

Supported Deployment Units

A *deployment unit* refers to a J2EE application (an Enterprise Application or Web application) or a stand-alone J2EE module (such as an EJB or Resource Adapter) that has been organized according to the J2EE specification and can be deployed to WebLogic Server.

For each type of deployment unit, the J2EE specification defines both the required files and their location in the directory structure of the application or module. Deployment units may include Java classes for EJBs and servlets, resource adapters, Web pages and supporting files, XML-formatted deployment descriptors, and even other modules.

J2EE does not specify *how* a deployment unit is deployed on the target server—only how standard applications and modules are organized. WebLogic Server supports the following types of deployment units:

- [“Enterprise Application” on page 2-7](#)
- [“Web Application” on page 2-7](#)
- [“Enterprise JavaBean” on page 2-8](#)

- “Resource Adapter” on page 2-8
- “Web Service” on page 2-8
- “J2EE Library” on page 2-8
- “Optional Package” on page 2-9
- “JDBC, JMS, and WLDF Modules” on page 2-9
- “Client Application Archive” on page 2-10

Enterprise Application

An Enterprise Application consists of one or more of the following J2EE applications or modules:

- Web applications
- Enterprise Java Beans (EJB) modules
- Resource Adapter modules

An Enterprise Application is packaged as an archive file with an `.ear` extension, but is generally deployed as an exploded EAR directory. An exploded EAR directory contains all of the JAR, WAR, and RAR modules (also in exploded format) for an application as well as the XML descriptor files for the Enterprise Application and its bundled applications and modules. See [Developing WebLogic Server Applications](#).

Web Application

A Web application always includes the following files:

- A servlet or JSP page, along with any helper classes.
- A `web.xml` deployment descriptor, a J2EE standard XML document that configures the contents of a WAR file.

Web applications may also contain JSP tag libraries, static `.html` and image files, supporting classes and `.jar` files, and a `weblogic.xml` deployment descriptor, which configures WebLogic Server-specific elements for Web applications. See [Developing Web Applications for WebLogic Server](#).

Enterprise JavaBean

Enterprise JavaBeans (EJBs) are reusable Java components that implement business logic and enable you to develop component-based distributed business applications. EJB modules are packaged as archive files having a `.jar` extension, but are generally deployed as exploded archive directories. The archive file or exploded archive directory for an EJB contains the compiled EJB classes, optional generated classes, and XML deployment descriptors for the EJB. See *Programming WebLogic Server Enterprise JavaBeans* for more information on the different types of EJBs.

Resource Adapter

A Resource Adapter (also referred to as a connector) adds Enterprise Information System (EIS) integration to the J2EE platform. Connectors provide a system-level software driver that WebLogic Server can use to connect to an EIS. Connectors contain both the Java classes, and if necessary, the native components required to interact with the EIS. See *Programming WebLogic Resource Adapters*.

Web Service

A Web Service is a set of functions packaged into a single entity that is available to other systems on a network, and can be shared by and used as a component of distributed Web-based applications. Web Services commonly interface with existing back-end applications, such as customer relationship management systems, order-processing systems, and so on.

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; typically the WAR or EJB JAR file is then packaged in an Enterprise Application. See *Programming WebLogic Web Services*.

J2EE Library

A J2EE library is a stand-alone J2EE module, or multiple J2EE modules packaged in an Enterprise Application (EAR), that is registered with the J2EE application container as a shared library at deployment time. After a J2EE library has been registered, you can deploy Enterprise Applications that reference the library in their `weblogic-application.xml` deployment descriptors. Each referencing application receives a copy of the shared J2EE library module(s) on deployment, and can use those modules as if they were packaged as part of the application itself. J2EE library support provides an easy way to share one or more J2EE modules among

multiple Enterprise Applications without physically adding the shared modules to each dependent application.

The deployment files of a shared library resemble either a standard Enterprise Application or J2EE module, as discussed in this section. Shared libraries differ from standard EARs and modules only by the contents of their `MANIFEST.MF` files. [Creating Shared J2EE Libraries and Optional Packages](#) in *Developing Applications with WebLogic Server* describes how to assemble and configure J2EE libraries, and how to configure Enterprise Applications that utilize J2EE libraries.

[“Deploying Shared Java EE Libraries and Dependent Applications” on page 6-22](#) describes how to deploy J2EE libraries and Enterprise Applications that reference J2EE libraries.

Optional Package

Optional packages provide similar functionality to J2EE libraries, allowing you to easily share a single JAR file among multiple applications. However, optional packages function as individual J2EE modules (stand-alone or within an Enterprise Application), rather than as an 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 are delivered as basic JAR files that have no deployment descriptors. You simply designate the JAR as an optional package at deployment time, and WebLogic Server registers the file with the target servers you select. After the optional package has been registered, you can then deploy J2EE modules and applications that reference the optional package in their `MANIFEST.MF` files. [Creating Shared J2EE Libraries and Optional Packages](#) in *Developing Applications with WebLogic Server* describes how to assemble and configure optional packages, and how to configure J2EE modules that utilize optional packages.

[“Deploying Shared Java EE Libraries and Dependent Applications” on page 6-22](#) describes how to deploy optional packages and J2EE modules that reference optional packages.

JDBC, JMS, and WLDF Modules

A JMS, JDBC, or WebLogic Diagnostic Framework (WLDF) application module can be deployed as a stand-alone resource, in which case the resource is available in the domain targeted during deployment, or as part of an Enterprise application. An application module deployed as part of an Enterprise Application is available only to the enclosing application (an *application-scoped resource*). Using application-scoped resources ensures that an application

always has access to required resources, and simplifies the process of deploying the application into new environments.

In contrast to system modules, application modules are owned by the developer who created and packaged the module, rather than the Administrator who deploys the module. This means that the Administrator has more limited control over JDBC, JMS, and WLDF application modules. When deploying an application module, an Administrator can change resource properties that were specified in the module, but cannot add or delete resources.

System modules are created by the Administrator via the WebLogic Administration Console, and can be changed or deleted as necessary by the Administrator. Similarly, stand-alone application modules created by the Administrator can be used to recreate global resources in multiple WebLogic Server environments simply by deploying the modules into new domains.

For more information on how to deploy and use JDBC, JMS, and WLDF modules, see:

- [“Deploying JDBC, JMS, and WLDF Application Modules” on page 6-11](#)
- [Configuring and Managing WebLogic JMS](#)
- [Configuring and Managing WebLogic JDBC](#)
- [Configuring and Using the WebLogic Diagnostic Framework](#)

Client Application Archive

The J2EE specification enables you to include a client application archive file within an Enterprise Application. A J2EE client application module contains the Java classes that execute in the client JVM (Java Virtual Machine) and deployment descriptors that describe EJBs (Enterprise JavaBeans) and other WebLogic Server resources used by the client. This enables both the server-side and client-side components to be distributed as a single unit. You define client modules in an EAR using the J2EE standard `application-client.xml` deployment descriptor and WebLogic Server `weblogic-applclient.xml` descriptor.

Deployment Tools

Weblogic Server provides the following tools to help you configure and deploy applications:

- [“weblogic.Deployer” on page 2-11](#)
- [“Administration Console” on page 2-11](#)
- [“WLST” on page 2-11](#)

- [“Deployment Tools for Developers” on page 2-11](#)

weblogic.Deployer

`weblogic.Deployer` provides a command-line based interface for performing both basic and advanced deployment tasks. Use `weblogic.Deployer` when you want command-line access to WebLogic Server deployment functionality, or when you need to perform a deployment task that is not supported using the Administration Console.

Administration Console

The Administration Console provides a series of Web-based Deployment Assistants that guide you through the deployment process. The Administration Console also provides controls for changing and monitoring the deployment status, and changing selected deployment descriptor values while the deployment unit is up and running.

Use the Administration Console when you need to perform basic deployment functions interactively and you have access to a supported browser.

WLST

The WebLogic Scripting Tool (WLST) is a new command-line interface that you can use to automate domain configuration tasks, including application deployment configuration and deployment operations. See [WebLogic Scripting Tool](#) for more information.

Deployment Tools for Developers

WebLogic Server provides several tools for deploying applications and stand-alone modules:

- `wldeploy` is an Ant task version of the `weblogic.Deployer` utility. You can automate deployment tasks by placing `wldeploy` commands in an Ant `build.xml` file and running Ant to execute the commands.
- `weblogic.PlanGenerator` is a command-line tools that enables developers to export an application’s configuration for deployment to multiple WebLogic Server environments.
- The deployment API allows you to perform deployment tasks programmatically using Java classes.
- The `autodeploy` domain directory allows you to deploy an application quickly for evaluation or testing in a development environment.

Understanding WebLogic Server Deployment

Preparing Applications and Modules for Deployment

The following sections provides information on key topics required to prepare applications for deployment:

- [“Packaging Files for Deployment” on page 3-2](#)
- [“Understanding Default Deployment Names” on page 3-4](#)
- [“Understanding Application Naming Requirements” on page 3-5](#)
- [“Understanding Deployment Version Strings” on page 3-6](#)
- [“Creating an Application Installation Directory” on page 3-6](#)
- [“Using FastSwap Deployment to Minimize Redeployment” on page 3-10](#)
- [“Best Practices for Preparing Deployment Files” on page 3-19](#)

Packaging Files for Deployment

WebLogic Server supports deployments that are packaged either as archive files using the `jar` utility or Ant's `jar` tool, or as exploded archive directories.

Note: In general, using archived files is more efficient when deploying applications to managed servers. However, it makes updating the application, such as updating web content, more difficult as it requires a redeployment of the application.

- [“Using Archived Files” on page 3-2](#)
- [“Using Exploded Archive Directories” on page 3-3](#)
- [“Creating an Exploded Archive Directory from an Archive File” on page 3-4](#)

Using Archived Files

An archive file is a single file that contains all of an application's or module's classes, static files, directories, and deployment descriptor files. In most production environments, the applications an Administrator receives for deployment are stored as archive files.

Deployment units that are packaged using the `jar` utility have a specific file extension depending on the type:

- EJBs and client archives are packaged as `.jar` files.
- Web applications are packaged as `.war` files.
- Resource adapters are packaged as `.rar` files.
- Enterprise applications are packaged as `.ear` files, and can contain other Java EE modules such as EJBs, JDBC, JMS, Web Applications, and Resource Adapters.
- Web Services can be packaged either as `.war` files or as `.jar` files, depending on whether they are implemented using Java classes or EJBs. Typically, the `.war` or `.jar` files are then packaged in an Enterprise Application `.ear` file.
- Java EE libraries are packaged either as an Enterprise Application (`.ear` file) or as a standard Java EE module.
- Client applications and optional packages are packaged as `.jar` files.

In addition to an archive file, you may also receive a deployment plan, which is a separate file that configures the application for a specific environment. [“Configuring Applications for Production Deployment” on page 4-1](#) describes deployment plans in more detail.

Using Exploded Archive Directories

An exploded archive directory contains the same files and directories as a JAR archive. If you choose to use an exploded archive directory, you may be required to manually unpack a previously-archived deployment. However, the files and directories reside directly in your file system and are not packaged into a single archive file with the `jar` utility.

You may choose to deploy from an exploded archive directory under the following circumstances:

- You want to perform partial updates of an Enterprise application after deployment. Deploying Enterprise Applications as an exploded archive directory makes it easier to update individual modules of the application without having to re-create the archive file.
- You are deploying a Web application or Enterprise application that contains static files that you will periodically update. In this case, it is more convenient to deploy the application as an exploded directory, because you can update and refresh the static files without re-creating the archive.
- You are deploying a Web application that performs direct file system I/O through the application context (for example, a Web application that tries to dynamically edit or update parts of the Web application itself). In this case, the modules that perform the I/O operations should have a physical filesystem directory in which to work; you cannot obtain a file when the application is deployed as an archive, as per the specification.

Java EE Rules for Deploying Exploded EAR Directories without Deployment Descriptors

The Java EE specification recommends that archived EARs (Enterprise Application Archives) can be deployed to a Java EE-compliant server without any deployment descriptors. To achieve this, all containers assume reasonable defaults *or* use annotated classes. In addition to supporting this mandate, WebLogic Server also allows deploying exploded EAR directories without deployment descriptors.

Since this applies to directories, certain rules are used to identify EARs and their nested modules. Otherwise, the WebLogic Server Administration Console or deployment tools will not treat the directories as valid exploded Java EE directories.

- For an exploded archived Web Application, in the absence of `WEB-INF/web.xml` descriptor, the name of the directory should have a `.war` suffix.
- For an exploded archived Enterprise Application without a `META-INF/application.xml` descriptor, the directory should have an `.ear` suffix. Within the application, the directory

of exploded Web module should have a `.war` suffix. Similarly, the exploded EJB module should have a `.jar` suffix and the exploded RAR module should have a `.rar` suffix.

- If an exploded Enterprise Application contains no `META-INF/application.xml` descriptor, the order in which modules are deployed is undefined and is dependent on the underlying `File.listFiles()` method order. To ensure a specific order in which modules are deployed, you must add an `application.xml` descriptor and list the modules in the desired order.

Creating an Exploded Archive Directory from an Archive File

If you have an archive file that you want to deploy as an exploded archive directory, use the `jar` utility to unpack the archive file in a dedicated directory. For example:

```
mkdir /myapp
cd /myapp
jar xvf /dist/myapp.ear
```

If you are unpacking an archive file that contains other module archive files (for example, an Enterprise Application or Web Service that includes JAR or WAR files) and you want to perform partial updates of those modules, you must expand the embedded archive files as well. Make sure that you unpack each module into a subdirectory having the same name as the archive file. For example, unpack a module named `myejb.jar` into a `/myejb.jar` subdirectory of the exploded Enterprise Application directory.

Note: If you want to use different subdirectory names for the archived modules in an exploded EAR file, you must modify any references to those modules in the application itself. For example, you must update the URI values specified in `application.xml` and `CLASSPATH` entries in the `manifest.mf` file.

Understanding Default Deployment Names

When you first deploy an application or stand-alone module to one or more WebLogic Server instances, you specify a deployment name to describe collectively the deployment files, target servers, and other configuration options you selected. You can later redeploy or stop the deployment unit on all target servers by simply using the deployment name. The deployment

name saves you the trouble of re-identifying the deployment files and target servers when you want to work with the deployment unit across servers in a domain.

If you do not specify a deployment name at deployment time, the deployment tool selects a default name based on the deployment source file(s). For archive files, `weblogic.Deployer` uses the name of the archive file without the file extension. For example, the file `myear.ear` has a default deployment name of `myear`. For an exploded archive directory, `weblogic.Deployer` uses the name of the top-level directory you deploy.

For Java EE libraries and optional packages, `weblogic.Deployer` uses the name specified in the library's manifest file. If no name was specified in the library's manifest file, you can specify one with the `-name` option.

See the following section, [“Understanding Application Naming Requirements”](#) for information on application naming requirements; See [“Deploying Applications and Modules with `weblogic.deployer`” on page 6-1](#) to specify a non-default deployment name.

Understanding Application Naming Requirements

In order to successfully deploy an application to WebLogic Server, the application name must be valid. Application naming requirements are as follows:

- Application names must only contain the following characters:
 - a-z
 - A-Z
 - 0-9
 - _ (underscore)
 - - (hyphen)
 - . (period)

No additional characters are allowed in application names.

- Application names that contain the “.” character must contain at least one additional different character; “. ” and “. . ” are not valid application names.
- Application names must be less than 215 characters in length.

Understanding Deployment Version Strings


In addition to a deployment name, an application or module can also have an associated version string. The version string distinguishes the initial deployment of the application from subsequent redeployed versions. For example, you may want to later update the application to fix problems or add new features. In production systems, it is critical to maintain a version string for both the initial and subsequent deployments of an application. Doing so allows you to update and redeploy an application version without interrupting service to existing clients. See [“Redeploying Applications in a Production Environment” on page 8-1](#) for more information.

The version string is specified in the manifest file for the application, and should be provided by your development team along with the other deployment files. [Assigning Application Versions](#) in *Developing WebLogic Server Applications* describes the conventions for specifying the version string.

Creating an Application Installation Directory

The application installation directory separates generated configuration files from the core application files, so that configuration files can be easily changed or replaced without disturbing the application itself. The directory structure also helps you to organize and maintain multiple versions of the same application deployment files.

The following figure shows the application installation directory hierarchy for storing a single version of a deployable application or module.

Figure 3-1 Application Installation Directory

Oracle recommends copying all new production deployments into an application installation directory before deploying to a WebLogic Server domain. Deploying from this directory structure helps you easily identify all of the files associated with a deployment unit—you simply deploy the installation root using the Administration Console, and the Console automatically locates associated files such as deployment plans and WebLogic Server deployment descriptors that were generated during configuration.

Steps for Creating an Application Installation Directory

To create an application installation directory:

1. Choose a top-level directory where you want to store deployment files for applications and modules on your system. Follow these best practices:
 - Do not store deployment files within a WebLogic Server domain directory.

Preparing Applications and Modules for Deployment

- Use source control if available to maintain deployment source files.
- If possible, store deployment files in a directory that is accessible by the Administration Server and Managed Servers in your domain.

The instructions that follow use the sample deployment directory,
`c:\deployments\production`.

2. Create a dedicated subdirectory for the application or module you want to deploy:

```
mkdir c:\deployments\production\myApplication
```

3. Create a subdirectory beneath the application directory to designate the version of the application you are deploying. Name the subdirectory using the exact version string of the application. For example:

```
mkdir c:\deployments\production\myApplication\91Beta
```

4. The version subdirectory will become the installation root directory from which you deploy the directory. Create subdirectories named `app` and `plan` under the version subdirectory:

```
mkdir c:\deployments\production\myApplication\91Beta\app
```

```
mkdir c:\deployments\production\myApplication\91Beta\plan
```

Note: If you have more than one deployment plan associated with the application, create one `\plan` subdirectory for each plan. For example, if you have two deployment plans associated with the `91Beta` version of the application `myApplication`, you would create two `\plan` subdirectories. For instance:

```
mkdir c:\deployments\production\myApplication\91Beta\plan1
```

```
mkdir c:\deployments\production\myApplication\91Beta\plan2
```

5. Copy your application source deployment files into the `\app` subdirectory. If you are deploying from an archive file, simply copy the archive file, as in:

```
cp c:\downloads\myApplication.ear  
c:\deployments\production\myApplication\91Beta\app
```

If you are deploying from an exploded archive directory, copy the complete exploded archive directory into `\app`:

```
cp -r c:\downloads\myApplication  
c:\deployments\production\myApplication\91Beta\app
```

This results in the new directory,

```
c:\deployments\production\myApplication\91Beta\app\myApplication.
```

6. If you have one or more deployment plans for the application, copy them into the `\plan` subdirectories.

If you have one deployment plan for the application:

```
cp c:\downloads\myApplicationPlans\plan.xml
c:\deployments\production\myApplication\91Beta\plan
```

If you have two deployment plans for the application:

```
cp c:\downloads\myApplicationPlans\plan1.xml
c:\deployments\production\myApplication\91Beta\plan1

cp c:\downloads\myApplicationPlans\plan2.xml
c:\deployments\production\myApplication\91Beta\plan2
```

Note: If you do not have an existing deployment plan, you can create one using the Administration Console as described in [“Configuring Applications for Production Deployment” on page 4-1](#). The Administration Console automatically stores newly-generated deployment plans in the `\plan` subdirectory of the application installation directory.

7. To install the application using Administration Console, select the application installation directory. By default, the Administration Console will use a plan named `plan.xml`, if one is available in the `\plan` subdirectory. The Administration Console does not identify plans in subdirectories other than the `\plan` subdirectory; in other words, plans in `\plan1` or `\plan2` subdirectories are not identified by the Administration Console. Therefore, if multiple plans for your application are available, you must indicate, in `config.xml`, the plan you would like to use. See [“Configuring Applications for Production Deployment” on page 4-1](#). For information on `config.xml`, see [Creating WebLogic Domains Using the Configuration Wizard](#).

After installing the application, you can configure, deploy, or distribute the application as necessary.

Note: You cannot specify an application installation directory when using the `weblogic.Deployer` tool, and the tool does not use an available `plan.xml` file by default. You must specify the actual deployment file(s) and plan to use for deployment. See [“Deploying Applications and Modules with weblogic.deployer” on page 6-1](#).

Using FastSwap Deployment to Minimize Redeployment

Today's web application developers expect to make changes to a deployed application and see those changes immediately by refreshing the browser. On the Java EE side, developers have to typically go through the following cycle to see their changes in action.

```
Edit -> Build -> Deploy -> Test
```

These steps, along with the many required descriptor elements, makes developing applications with Java EE seem complex and top-heavy. Among these steps, the Build and Deploy cycles are necessitated by Java and by the application server being employed. IDEs are trying to make the Edit and Build steps seamless by providing incremental compilation support. On the server side, the WLS FastSwap deployment feature makes the Deploy and Test cycles just as seamless.

How FastSwap Deployment Works

Java EE 5 introduces the ability to redefine a class at runtime without dropping its `ClassLoader` or abandoning existing instances. This allows containers to reload altered classes without disturbing running applications, vastly speeding up iterative development cycles and improving the overall development and testing experiences. The usefulness of the Java EE dynamic class redefinition is severely curtailed, however, by the restriction that the shape of the class – its declared fields and methods – cannot change. The purpose of FastSwap is to remove this restriction in WLS, allowing the dynamic redefinition of classes with new shapes to facilitate iterative development.

With FastSwap, Java classes are redefined in-place without reloading the `ClassLoader`, thereby having the decided advantage of fast turnaround times. This means that you do not have to wait for an application to redeploy and then navigate back to wherever you were in the Web page flow. Instead, you can make your changes, auto compile, and then see the effects immediately.

Supported FastSwap Application Configurations

The following application configuration are supported when using FastSwap deployment:

- FastSwap is only supported when WLS is running in development mode. It is automatically disabled in production mode.
- Only changes to class files in exploded directories are supported. Modifications to class files in archived applications, as well as archived JAR files appearing in the application's classpath are not supported. Examples are as follows:

- When a web application is deployed as an archived WAR within an EAR, modifications to any of the classes are not picked up by the FastSwap agent.
- Within an exploded web application, modifications to Java classes are only supported in the `WEB-INF/classes` directory; the FastSwap agent does not pick up changes to archived jars residing in `WEB-INF/lib`.

Enabling FastSwap In Your Application

To enable FastSwap in your application, add the following element to the `weblogic-application.xml` file.

```
<fast-swap>true</fast-swap>
```

For more information on the `weblogic-application.xml` elements, see [Enterprise Application Deployment Descriptor Elements](#) in *Developing Applications with WebLogic Server*.

FastSwap can also be enabled for a standalone web-application by adding the `<fast-swap>` element to the `weblogic.xml` file. For more information on the `weblogic.xml` elements, see [weblogic.xml Deployment Descriptor Elements](#) in *Developing Web Applications, Servlets, and JSPs for WebLogic Server*.

Overview of the FastSwap Process

The following steps describe how the FastSwap deployment process works:

1. Once FastSwap is enabled at the descriptor level, an appropriate `ClassLoader` is instantiated when the application is deployed to WLS.
2. Open a browser to see the application at work. Then modify (add/edit/delete) the methods and/or classes (see [Limitations When Using FastSwap](#)) and compile them.

It is recommended that you use an IDE such as Eclipse or IntelliJ and setting the compile-on-save option so that java files are compiled on saving. Also note that the FastSwap agent does not compile Java files.

3. Refresh the browser or send a new request to the application.

The FastSwap agent tries to find all classes that have been modified since the last iteration by looking at all directories in the classpath. Considering an exploded application with a single web application, the following directories are examined for any class file modifications based on their timestamps:

```
ExampleApp/APP-INF/classes
ExampleApp/webapp/WEB-INF/classes
```

The FastSwap agent redefines the modified classes in the application and then serves the request.

Using Ant with the JMX Interface

For *headless* applications (that is, applications not fronted by a web application), class redefinition can be explicitly initiated using the JMX interface. An Ant task that uses the JMX interface can be used to initiate class redefinition, as shown in following Ant FastSwapTask example.

Listing 3-1

```
<project name='MyProject' default='all' >
  <taskdef name='fast-swap'
classname='com.bea.wls.redef.ant.FastSwapTask' />
  <target name='all'>
    <!--
      Redefine classes which have changed since they were last loaded.
      Required parameters:
        adminUrl: Connection url
        user: User name
        password: User password
        server: Managed server name
        application: Deployed application name
      Optional parameters:
        module: Name of the module within the application.
                If not specified, all modules within the application
                will be processed.
        failonerror: Default=true. If true, task will fail if fast-swap failed.
                    Otherwise, a message will be displayed and the task will
                    return success.
        timeout: Default=300. Timeout in seconds.
        classnames: Comma separated list of classnames to process. If not
                    specified, all classes within specified modules (if any)
                    in the application will be considered.
    -->
  </target>
</project>
```

```

-->
<fast-swap
  adminUrl='t3://localhost:7001'
  user='weblogic'
  password='weblogic'
  server='myserver'
  application='SimpleApp'

  module='SimpleAppCookie'
  failonerror='false'
  timeout='30'
  classnames='examples.servlets.CookieCounter1,
              examples.servlets.CookieCounter2,
              examples.servlets.CookieCounter'

/>
</target>
</project>

```

Application Types and Changes Supported with FastSwap

FastSwap is supported with POJOs (JARs), Web applications (WARs) and enterprise applications (EARs) deployed in an exploded format. FastSwap is not supported with resource adapters (RARs).

The following types of changes are supported with FastSwap:

- Addition of static methods
- Removal of static methods
- Addition of instance methods
- Removal of instance methods
- Changes to static method bodies
- Changes to instance method bodies
- Addition of static fields

- Removal of static fields
- Addition of instance fields
- Removal of instance fields

The following table lists detailed change types supported with FastSwap:

Table 3-1 Supported Application Types and Changes

Scope	Java Change Type	Supported	Notes
Java Class	Add method	Yes	Addition of the <code>finalize</code> method is not supported.
Instance (non-abstract)	Remove method	Yes	Addition of the <code>finalize</code> method is not supported.
	a) Add field	Yes	
	b) Remove field	Yes	
	c) Change method body	Yes	
	d) Add constructor	Yes	
	e) Remove constructor	Yes	
	f) Change field modifiers	Yes	
	g) Change method modifiers	Yes	
	Class-level (static)	Add method	Yes
	Remove method	Yes	
	Change body method	Yes	

Table 3-1 Supported Application Types and Changes

Scope	Java Change Type	Supported	Notes
Class Hierarchy Changes	Change list of implemented interfaces	No	
	Change extends “SuperClass”	No	
Abstract Java Class	Add abstract method	Yes	
	Delete abstract method	Yes	
	All other supported changes (a–g) listed in Instance	Yes	
“final” Java Class	Same supported changes (a–g) listed in Instance	Yes	
“final” Java Method	Same supported changes (a–g) listed in Instance	Yes	
“final” Java Field	Same supported changes (a–g) listed in Instance	Yes	
Enum	Add constants	No	
	Remove constants	No	
	Add/remove methods	No	
Anonymous Inner Class	Add/remove fields	NA	Not supported by the Java language
	Add/remove methods	No	

Table 3-1 Supported Application Types and Changes

Scope	Java Change Type	Supported	Notes
Static Inner Class	Same supported changes (a–g) listed in Instance	Yes	
Member Inner Classes (non-static inner classes)	Same supported changes (a–g) listed in Instance	Yes	
Local Inner Classes	Same supported changes (a–g) listed in Instance	Yes	
Java Interface	Add method	Yes	
Java Reflection	Access existing fields/methods	Yes	
	Access new methods	No	New methods are not seen using Reflection and some synthetic methods are exposed.
	Access new fields	No	New fields are not seen using Reflection.
Annotations on Classes	Add or remove method/field annotations	No	
Annotation Type	Add or remove methods/attributes	No	
Exception Classes	Same supported changes (a–g) listed in Instance	Yes	
EJB Interface	Add/remove methods	No	Changes to EJB interfaces involve Reflection, which is not fully supported.
EJB 3.0 Session/MDB	Add/remove methods	No	Any support classes referenced by the EJB classes can be modified.
EJB Implementation Class	Add/remove fields	No	

Table 3-1 Supported Application Types and Changes

Scope	Java Change Type	Supported	Notes
EJB 3.0 EntityBean	Add/remove methods	No	Any support classes referenced by the EJB classes can be modified.
	Add/remove fields	No	
EJB Interceptors	Add/remove methods	No	Any support classes referenced by the EJB classes can be modified.
	Add/remove fields	No	

Limitations When Using FastSwap

The following limitation apply when using FastSwap deployment:

- Java Reflection results do not include newly added fields and methods and include removed fields and methods. As a result, use of the reflection API on the modified classes can result in undesired behavior.
- Changing the hierarchy of an already existing class is not supported by FastSwap. For example, either a) changing the list of implemented interfaces of a class; or b) changing the superclass of a class, is not supported.
- Addition or removal of Java annotations is not supported by FastSwap, since this is tied to the reflection changes mentioned above.
- Addition or removal of methods on EJB Interfaces is not supported by FastSwap since an EJB Compilation step is required to reflect the changes at runtime.
- Addition or removal of constants from Enums is not supported.
- Addition or removal of the finalize method is not supported.

Handling Unsupported FastSwap Changes

When FastSwap is enabled, after you recompile a class, FastSwap attempts to redefine classes in existing classloaders. If redefinition fails because your changes fall outside the scope of supported FastSwap changes, the JVM throws an `UnsupportedOperationException` in the WLS window and in the server log file. Your application will not reflect the changes, but will continue to run.

Preparing Applications and Modules for Deployment

To implement your changes, you can redeploy the application or affected modules (partial redeploy), depending on the application type and the extent of your changes.

Best Practices for Preparing Deployment Files

The following best practices are recommended when preparing applications and modules for deployment:

- Regardless of whether you deploy an archive file or exploded archive directory, store the deployment files in an installation directory for the application, as described in [“Creating an Application Installation Directory”](#) on page 3-6. Using an installation directory simplifies the deployment process, because the Administration Console understands where to locate deployment and configuration files.
- Manage the entire application installation directory in a source control system, so you can easily revert to previous application versions if necessary.

Preparing Applications and Modules for Deployment

Configuring Applications for Production Deployment

The following sections describe how you configure applications for deployment to a production WebLogic Server environment:

- [“Understanding the Deployment Configuration Process”](#) on page 4-1
- [“Typical Deployment Configuration Workflows”](#) on page 4-6
- [“Creating a New Deployment Plan to Configure an Application”](#) on page 4-11
- [“Understanding Deployment Plan Contents”](#) on page 4-13
- [“Using an Existing Deployment Plan to Configure an Application”](#) on page 4-17
- [“Generic File Loading Overrides”](#) on page 4-19
- [“Additional Configuration Tasks”](#) on page 4-21
- [“Best Practices for Managing Application Configuration”](#) on page 4-21

Understanding the Deployment Configuration Process

When an administrator or deployer receives a new application, or a new version of an application, from development or quality assurance teams, the application is usually configured for a development or testing environment. The application may use specific resource names and performance tuning settings that match the available resources on the target servers used in the development or QA environments where the application was last deployed.

Because development and testing environments can be significantly different from the production environment in which the application is ultimately deployed, an Administrator must configure the application to use resource names and performance tuning parameters that are valid and appropriate for the production environment.

- [“Deployment Configuration Lifecycle” on page 4-2](#)
- [“Understanding Application Deployment Descriptors” on page 4-3](#)
- [“Understanding WebLogic Server Deployment Plans” on page 4-4](#)
- [“Goals for Production Deployment Configuration” on page 4-5](#)

Deployment Configuration Lifecycle

Deployment configuration for an application can occur at several points in the lifecycle of an application. Each phase of deployment configuration typically involves creating and working with different deployment files:

1. Development configuration—During development, a programmer creates J2EE deployment descriptors for an application or module. The programmer also creates WebLogic Server deployment descriptors to configure the application for deployment to a WebLogic Server development environment. See [Developing Applications with WebLogic Server](#).

Note: Applications developed outside of the WebLogic Server development environment (for example, a sample or third-party J2EE application such as PetStore) may include only J2EE descriptors.

2. Export configuration—Before releasing an application from development, a programmer or designer may optionally export the application’s deployment configuration to a WebLogic Server deployment plan. Exporting a configuration creates deployment plan variables for all or a subset of the deployment properties already defined by a developer in the application’s WebLogic Server deployment descriptor files. See [“Exporting an Application for Deployment to New Environments” on page 5-1](#).

Exporting an application helps deployers in other areas of the organization (such as engineers on the QA team or production Administrators) easily deploy the application to environments that differ from the programmer’s development environment. The ideal deployment plan includes all of the properties that a deployer needs to change before deploying the application in a new environment.

3. Deployment-time configuration—An Administrator or deployer configures the application before deploying the application into the target environment. Deployment-time configuration may use:

- The same WebLogic Server deployment configuration and deployment plan created during development
- Modified versions of the development configuration and deployment plan
- A custom deployment plan that the deployer previously created for the environment, depending on the deployment configuration workflow for your organization.

See [“Deploying Applications and Modules with weblogic.deployer”](#) on page 6-1.

4. Post-deployment configuration—After an application has been deployed to a target environment, an Administrator or deployer can reconfigure the application by redeploying with a new deployment plan or by using the Administration Console to update and redeploy an existing deployment plan. See [“Redeploying Applications in a Production Environment”](#) on page 8-1 and [“Managing Deployed Applications”](#) on page 9-1.

Because deployment configuration is performed by different people at different points in the lifecycle of an application, administrators, deployers, and developers need to work together to create a repeatable configuration workflow for their organization. See [“Typical Deployment Configuration Workflows”](#) on page 4-6.

Understanding Application Deployment Descriptors

The basic deployment configuration for an application is defined in multiple XML documents, known as deployment descriptors, that are included as part of the application archive file that you receive for deployment. Deployment descriptor files fall into two separate categories:

- *J2EE deployment descriptors* define the fundamental organization and behavior of a J2EE application or module, regardless of where the application is deployed. Each J2EE application and module requires a specific J2EE deployment descriptor as defined in the J2EE 1.4 specification.
- *WebLogic Server deployment descriptors* define the resource dependencies and tuning parameters that an application uses in a specific WebLogic Server environment.

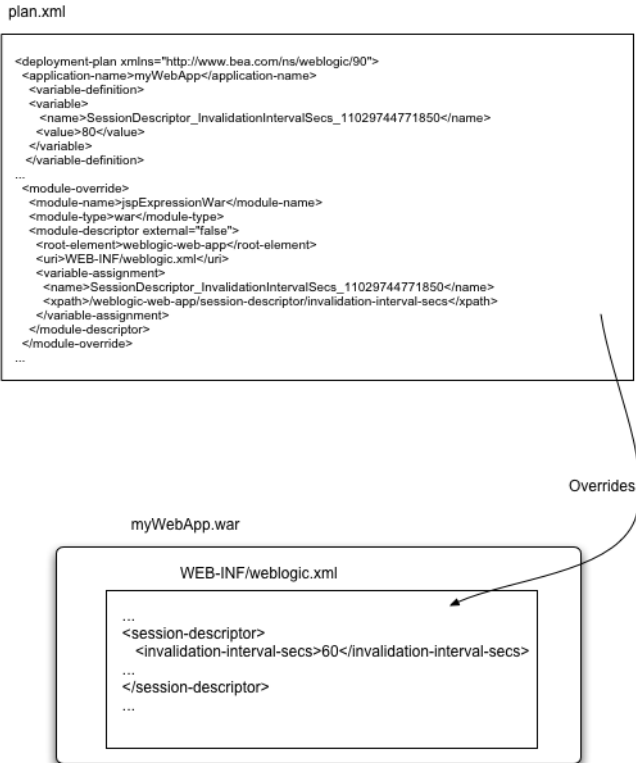
For the purposes of a production deployment, you should treat both the J2EE and WebLogic Server deployment descriptors as part of the application’s source code, which is owned by your development team. Do not edit application deployment descriptors in order to configure an application for deployment to a production environment. Instead, persist configuration changes into a WebLogic Server *deployment plan*, which is described in the next section.

Understanding WebLogic Server Deployment Plans

As previously discussed, a WebLogic Server deployment plan is an optional XML document that you use to configure an application for deployment to a specific WebLogic Server environment. A deployment plan specifies setting deployment property values that would normally be defined in an application's WebLogic Server deployment descriptors, or overrides property values that are already defined in a WebLogic Server deployment descriptor. When exporting an application, the deployment plan typically acts to override selected properties in the WebLogic Server deployment descriptors you created during development.

Typically, deployment plans are created by developers along with the associated application files, then distributed to the Administrator or another deployer, who updates the plan for a particular environment (such as staging, testing, or production). The deployment plan is stored *outside* of an application archive or exploded archive directory. As a best practice, Oracle recommends storing each deployment plan for a single application in its own `plan` subdirectory of the application's root directory (See [“Creating an Application Installation Directory”](#) on page 3-6).

Deployment plans help the Administrator easily modify an application's WebLogic Server configuration for deployment to multiple, differing WebLogic Server environments *without* modifying the deployment descriptor files included in the application archive. Configuration changes are applied by adding or changing variables in the deployment plan, which define both the location of the WebLogic Server descriptor properties to change and the value to assign to those properties. Administrators deploying an application need only change the deployment plan—the original deployment files and deployment descriptors remain unchanged as shown in [Figure 4-1](#). To determine the deployment configuration workflow for your environment, see [“Typical Deployment Configuration Workflows”](#) on page 4-6.

Figure 4-1 WebLogic Server Deployment Plan

Goals for Production Deployment Configuration

For the Administrator, the primary goal of configuring an application for production deployment is to generate a new deployment plan that is valid and appropriate for the target WebLogic Server environment. Specifically, the deployment plan must resolve all external resources references for the application to refer to valid resources available in the target environment. If the Application's configuration does not define to valid external resources for the target servers, the application cannot be deployed.

A deployment plan can optionally define or override WebLogic Server tuning parameters, to make ideal use of resources in the target environment. Defining tuning parameters is not required

in order to successfully deploy an application. If an application's deployment descriptors and deployment plan do not define tuning parameters, WebLogic Server uses default values.

Typical Deployment Configuration Workflows

Deployment plans enable you to define a convenient, repeatable workflow for configuring an application for deployment to multiple WebLogic Server environments. A configuration workflow for production applications requires cooperation between your development and design teams, which create and package the deployable application, and the administrator or deployer for each target WebLogic Server environment.

The ideal deployment configuration workflow for your organization is determined by:

- Number of different environments in which you deploy the same application
- Difference in resources provided by each target environment
- Frequency with which each target environment changes
- Frequency with which the application's J2EE configuration changes
- Ownership requirements for configuration information in different areas of your organization

The sections that follow describe common deployment configuration workflows for managing deployment plans and deploying applications to multiple WebLogic Server domains.

Caution: Oracle does not support using a deployment plan to change the `context-root` in an `application.xml` file. However, if an application is deployed as a library, you can either change the `context-root` through an `weblogic-application.xml` file or use the deployment plan to change the `context-root` in an `weblogic-application.xml` file.

Application with Single Deployment Plan

Organizations that know the exact configuration of different deployment environments can use a single, well-defined deployment plan to deploy an application to multiple WebLogic Server domains. The single deployment plan configuration workflow works in the following way:

1. The development team, cooperating with administrators and deployers, creates a master deployment plan for use with all target environments. The number of target environments will vary depending on your organizational structure. Common deployment environments include

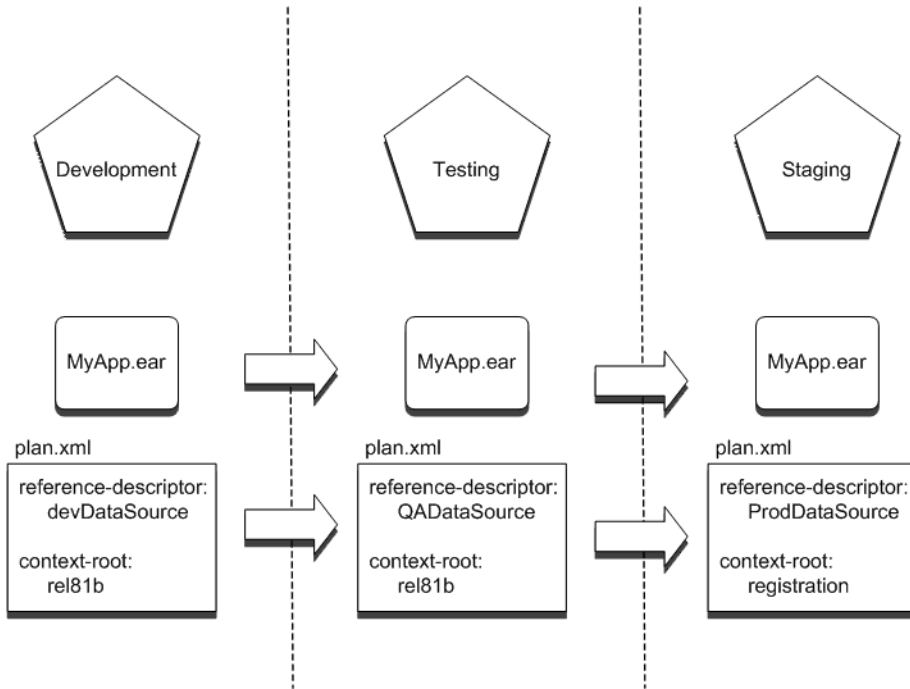
one or more Quality Assurance (QA) or testing domains, staging domains, and production domains.

The deployment plan that the team creates at this phase defines variables for all configuration properties that are known to differ between each target environment. For example, the plan might define empty variables for resource names that differ between environments and must be configured before the application can be deployed. The plan may also define default values for common tuning parameters that deployers may want to change in their environments.

For more information about creating a deployment plan during development, see [“Exporting an Application for Deployment to New Environments” on page 5-1](#).

2. When a version of the application is ready to be released, the development team packages the application deployment files and delivers both the deployment files and a master deployment plan to deployers for each target environment.
3. Each deployer uses the Administration Console to install the application and identify the deployment plan to use for configuration. The Administration Console validates the overall deployment configuration based on the resources available in the target domain. The Console then presents a list of configurable properties defined in the plan (as well as any invalid properties) to the deployer for editing.

Figure 4-2 Single Deployment Plan Workflow



4. The deployer uses the Administration Console to interactively configure properties that were defined in the deployment plan. Deployment plan variables that have null values, or invalid values for the target WebLogic Server instances or clusters, must be configured before the application can be deployed. Deployment plan variables that already have valid values need not be changed before deployment.

Deployers in each environment agree to limit their configuration changes to those properties defined in the deployment plan. If additional configuration changes are required, the deployer must communicate those requirements to the development or design team that modifies the master deployment plan.

Benefits of a Single Deployment Plan Workflow

Using the single deployment plan workflow provides the following benefits:

- It enables the development or design team to control the deployment configuration of the application.

- It reduces the number of configuration decisions that a deployer must make when deploying the application to a target environment.
- It minimizes the number of configuration files associated with the application, making it easy to store deployment configuration information in a source control system.

In general, you would use a single deployment plan workflow if your organization has a few, well-understood target environments, and you want to easily replicate a standardized deployment configuration in each environment.

Single Deployment Plan Ownership and Limitations

The single deployment plan workflow assumes that the development or design team maintains ownership of the deployment plan, and that deployers limit their plan changes to those variables defined in the plan. If the deployer modifies only those properties defined in the deployment plan, their changes are written back to the same deployment plan as updates to the variables.

However, WebLogic Server imposes no restrictions on the configuration properties that a deployer can modify using the Administration Console. If a deployer configures deployment properties that were not originally defined in a plan, the Console generates a new deployment plan having the additional variable entries, and uses the new plan for deployment or redeployment operations. This can lead to a situation where the deployer uses a deployment plan that is drastically different from the master deployment plan owned by the development team.

To incorporate new changes into the master deployment plan, the deployer retrieves the new, customized deployment plan created by the Console. Ideally, those changes should be applied to the master deployment plan.

Application with Multiple Deployment Plans

Organizations that have numerous deployment environments that frequently change should use a configuration workflow with multiple deployment plans. In a multiple deployment plan workflow, each deployment plan is owned by the deployer of the application rather than the development team. The multiple deployment plan configuration workflow works in the following way:

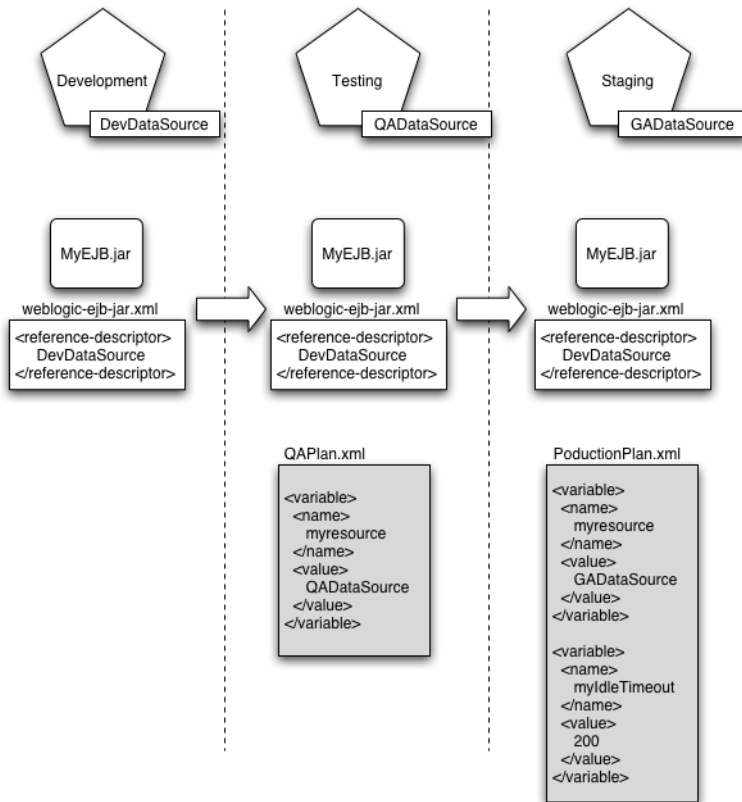
1. The development team releases a version of the packaged application deployment files (containing J2EE and Weblogic Server descriptors). The development team may or may not include a template deployment plan with exported variables for resource definitions or common tunable parameters.

- Before deploying the application, each deployer generates a custom deployment plan to configure the application for their target environment.

A custom deployment plan can be created by starting with a template deployment plan (or no deployment plan) and making changes to the application’s deployment configuration using the Administration Console. See [Update a deployment plan](#) in *Administration Console Online Help*.

- After defining the deployment configuration for their environment, each deployer retrieves their custom deployment plan and maintains it for future deployments of the application. Oracle recommends storing custom configuration plans in a source control system so that new versions can be tracked and reverted to if necessary.

Figure 4-3 Multiple Deployment Plan Workflow



4. For subsequent releases of the application, each deployer uses their customized deployment plan to configure the application for deployment. Using the customized plan allows deployers to perform deployments with the `weblogic.Deployer` or automate deployments using WLST.

Benefits of a Multiple Deployment Plan Workflow

Using the multiple deployment plan workflow provides the following benefits:

- It enables the administrator or deployer to manage both the application configuration and environment configuration in tandem.
- It enables deployers to automate the deployment process by using a custom plan that fully configures the application for their application.

In general, you would use a multiple deployment plan workflow if your organization has many deployment environments that change frequently, making it difficult or impossible to maintain a single master deployment plan.

Multiple Deployment Plan Ownership and Limitations

The multiple deployment plan workflow assumes that the deployer or administrator (rather than the programming or design team) owns and maintains the deployment configuration for an application. It also assumes that the basic J2EE configuration of the application rarely changes, because certain J2EE configuration changes would render a deployer's custom configuration plans invalid. For example, if a module in an Enterprise Application is added, removed, or changed, custom deployment plans referencing the module would become invalid. In this case, each deployer would need to re-create a custom plan by configuring the application using the Administration Console.

Creating a New Deployment Plan to Configure an Application

The Administration Console automatically generates (or updates) a valid XML deployment plan for an application when you interactively change deployment properties for an application that you have installed to the domain. You can use the generated deployment plan to configure the application in subsequent deployments, or you can generate new versions of the deployment plan by repeatedly editing and saving deployment properties.

Notes:

- `weblogic.PlanGenerator` also enables you to generate a basic WebLogic Server deployment plan for applications that have only J2EE deployment descriptors. See [“weblogic.PlanGenerator Command Line Reference” on page B-1](#).
- You cannot use a deployment plan to change the `context-root` in an `application.xml` file. However, if an application is deployed as a library, you can either change the `context-root` through an `weblogic-application.xml` file or use the deployment plan to change the `context-root` in an `weblogic-application.xml` file.

Generating a deployment plan using the Administration Console involves these steps:

1. [“Preparing the Deployment Files” on page 4-12](#)
2. [“Installing the Application Archive” on page 4-12](#)
3. [“Saving Configuration Changes to a Deployment Plan” on page 4-13](#)

The sections that follow use the WebLogic Server sample application `jspExpressionEar.ear` to describe each step.

Preparing the Deployment Files

Use the following the directions to prepare your application for deployment.

1. Create a new root directory and `app` and `plan` subdirectories for your application. For example:

```
mkdir c:\sample_root
mkdir c:\sample_root\app
mkdir c:\sample_root\plan
```

2. Download the WebLogic Server sample application `jspExpressionEar.ear` from <http://e-docs.bea.com/wls/docs103/jspExpressionEar.ear>. Be sure to save the application to the `c:\sample_root\app` directory you created in Step 1.

Installing the Application Archive

The Administration Console uses an application installation assistant to help you install a new application for configuration and deployment to a new WebLogic Server environment. After you install an application or module and select deployment targets, the deployment files are available in the WebLogic Server domain and can be configured, distributed, and deployed as necessary.

Follow these steps to install the sample application to the examples server domain:

1. Start the examples server WebLogic server by using the Windows start menu or by running the `WL_HOME\samples\domains\wl_server\startWebLogic.cmd` script.
2. Access the Administration Console by pointing your browser to `http://localhost:7001/console`.
3. Log in to the Administration Console.
4. Follow the steps in [Install applications and modules](#) in *Administration Console Online Help* to install your application or the `jspExpressionEar.ear` sample application you downloaded in [“Preparing the Deployment Files”](#) on page 4-12.

Saving Configuration Changes to a Deployment Plan

Use the Administration Console to edit deployment configuration properties for the application you installed in [“Installing the Application Archive”](#) on page 4-12 and save the configuration to a deployment plan. For example, you could change properties such as the following in the `jspExpressionEar.ear` sample application:

1. On the **Configuration** page, edit one or more configuration properties. For example, change the Session Invalidation Interval to 80 seconds, and the Session Timeout to 8000 seconds.
2. Click **Save** to save your changes. The Administration Console stores your configuration changes to a new deployment plan. If you deployed the sample application from a root directory, the Administration Console automatically places the new deployment plan in the `\plan` subdirectory of the root directory. For example, `c:\sample_root\plan\Plan.xml`.

Understanding Deployment Plan Contents

The deployment plan generated in [“Creating a New Deployment Plan to Configure an Application”](#) on page 4-11 contains the entries shown in [“Sample Deployment Plan”](#) on page 4-13.

Listing 4-1 Sample Deployment Plan

```
<deployment-plan xmlns="http://www.bea.com/ns/weblogic/90">
  <application-name>sample_root</application-name>
  <variable-definition>
    <variable>
      <name>SessionDescriptor_InvalidationIntervalSecs_11029744771850</nam
e>
```

Configuring Applications for Production Deployment

```
    <value>80</value>
  </variable>
  <variable>
    <name>SessionDescriptor_TimeoutSecs_11029744772011</name>
    <value>8000</value>
  </variable>
</variable-definition>
<module-override>
  <module-name>jspExpressionEar.ear</module-name>
  <module-type>ear</module-type>
  <module-descriptor external="false">
    <root-element>weblogic-application</root-element>
    <uri>META-INF/weblogic-application.xml</uri>
  </module-descriptor>
  <module-descriptor external="false">
    <root-element>application</root-element>
    <uri>META-INF/application.xml</uri>
  </module-descriptor>
</module-override>
<module-override>
  <module-name>jspExpressionWar</module-name>
  <module-type>war</module-type>
  <module-descriptor external="false">
    <root-element>weblogic-web-app</root-element>
    <uri>WEB-INF/weblogic.xml</uri>
    <variable-assignment>
      <name>SessionDescriptor_InvalidationIntervalSecs_11029744771850</n
ame>
      <xpath>/weblogic-web-app/session-descriptor/invalidation-interval-
secs</xpath>
    </variable-assignment>
    <variable-assignment>
      <name>SessionDescriptor_TimeoutSecs_11029744772011</name>
      <xpath>/weblogic-web-app/session-descriptor/timeout-secs</xpath>
    </variable-assignment>
  </module-descriptor>
  <module-descriptor external="false">
    <root-element>web-app</root-element>
```



```

    <uri>WEB-INF/web.xml</uri>
  </module-descriptor>
</module-override>
<module-override>
  <module-name>sample_root</module-name>
  <module-type>ear</module-type>
  <module-descriptor external="false">
    <root-element>weblogic-application</root-element>
    <uri>META-INF/weblogic-application.xml</uri>
  </module-descriptor>
  <module-descriptor external="false">
    <root-element>application</root-element>
    <uri>META-INF/application.xml</uri>
  </module-descriptor>
</module-override>
<config-root>C:\sample_root\plan</config-root>
</deployment-plan>

```

The basic elements in the deployment plan serve the following functions:

- `deployment-plan` encapsulates all of the deployment plan's contents.
- `application-name` corresponds to the deployment name for the application or module.
- `variable-definition` defines one or more `variable` elements. Each `variable` defines the name of a variable used in a plan and a `value` to assign (which can be null). The sample plan shown in [“Sample Deployment Plan” on page 4-13](#) contains variable definitions for the changes you made to the Session Invalidation Interval and Session Timeout properties.
- `module-override` elements define each module name, type, and deployment descriptor that the deployment plan overrides. A `module-descriptor` element can optionally contain a `variable-assignment` which identifies a variable name used to override a property in the descriptor, and the exact location within the descriptor where the property is overridden.

The sample plan shown in [“Sample Deployment Plan” on page 4-13](#) contains module override elements for the Enterprise Application, the embedded Web application, and the enclosing root directory. The `module-descriptor` entry for the `weblogic.xml` descriptor file

contains two `variable-assignment` elements that override the property values for the Session Invalidation Interval and Session Timeout properties you changed in [“Saving Configuration Changes to a Deployment Plan”](#) on page 4-13.

By default, the values in `variable-assignment` elements are added to the values that are already defined in the descriptor. You can change this behavior and cause the `variable-assignment` element to replace or remove the values that are defined in the descriptor by setting the `operation` sub-element in the `variable-assignment` element to the value `replace` or `remove`, respectively.

For example, suppose that in `ejb-jar.xml`, a developer created a policy to allow access only to the security role named `ejbRole`.

```
...
<assembly-descriptor>
  <security-role>
    <role-name>ejbRole</role-name>
  </security-role>
  <method-permission>
    <role-name>ejbRole</role-name>
    <method>
      <ejb-name>ejb.SearchHandlerWrapperEJB</ejb-name>
      <method-name>*</method-name>
    </method>
  </method-permission>
</assembly-descriptor>
...
```

And, the `security-role-assignment` element in `weblogic-ejb-jar.xml`, `ejbRole` is mapped to the principal named `user1`.

```
...
<security-role-assignment>
  <role-name>ejbRole</role-name>
  <principal-name>user1</principal-name>
</security-role-assignment>
...
```

If you want to use a deployment plan to override the `security-role-assignment` element defined in `weblogic-ejb-jar.xml`, so that `ejbRole` is mapped to `user2` instead of `user1`, you could achieve the desired override behavior by setting appropriate values for the `variable`, `variable-assignment`, and `operation` elements in the deployment plan. Make sure to set the value of `operation` to `replace`:

```
...
<variable>
  <name>SecurityRoleAssignment_ejbRole_PrincipalNames_11168815313911</
  name>
```

```

<value>user2</value>
</variable>

<variable-assignment>
<name>SecurityRoleAssignment_ejbRole_PrincipalNames_11168815313911</
name>
<xpath>/weblogic-ejb-jar/security-role-assignment/[role-name="ejbRol
e"]/principal-name</xpath>
<operation>replace</operation>
</variable-assignment>

```

For more information about the contents of a WebLogic Server deployment plan, see <http://www.bea.com/ns/weblogic/deployment-plan/1.0/deployment-plan.xsd>.

Using an Existing Deployment Plan to Configure an Application

Applications that you receive for deployment may come with varying levels of configuration information. If you have an existing deployment plan for an application, simply prepare the application as described in [“Preparing the Deployment Files” on page 4-12](#) and place the deployment plan in the `plan` subdirectory of the application root. Then install the application using the instructions in [“Installing the Application Archive” on page 4-12](#). The Administration Console automatically uses a deployment plan named `plan.xml` in the `\plan` subdirectory of an application root directory if one is available. If multiple plans are available for your application, they are placed in their own `\plan` subdirectories (for example `\plan1` and `\plan2`), and the Administration Console cannot identify them. Therefore, the `config.xml` must specify the plan you want to use. For information on `config.xml`, see [Domain Configuration Files in Understanding Domain Configuration](#).

After you install a new application and existing deployment plan, the Administration Console validates the deployment plan configuration against the target servers and clusters that were selected during installation. If the deployment plan contains empty (null) variables, or if any values configured in the deployment plan are not valid for the target server instances, you must override the deployment plan before you can deploy the application. You can also configure tuning parameters to better suit the target environment in which you are deploying the application, as described in [“Saving Configuration Changes to a Deployment Plan” on page 4-13](#). Changes you make to the application’s configuration are saved to a new deployment plan.

If you have a valid deployment plan that fully configures an application for the environment in which you are deploying, you can use either the Administration Console or the `weblogic.Deployer` utility to deploy an application with a deployment plan to use for deployment.

Note: Any deployment plan you use with the `weblogic.Deployer` utility must be complete and valid for your target servers. `weblogic.PlanGenerator` does not allow you to set or override individual deployment properties when it creates a plan. To deploy a new application and existing deployment plan using `weblogic.Deployer`, see [“Deploying an Application with a Deployment Plan” on page 6-4](#).

Generic File Loading Overrides

This feature allows you to place application-specific files to be overridden into a new optional subdirectory (named `/AppFileOverrides`) in the existing plan directory structure. The presence or absence of this new optional subdirectory controls whether file overrides are enabled for the deployment. If this subdirectory is present, an internal `ClassFinder` is added to the front of the application and module `ClassLoaders` for the deployment. As a result, the file override hierarchy rules follow the existing `ClassLoader` and resource loading rules and behaviors for applications. For more information on WLS application classloading, see [WebLogic Server Application Classloading](#) in *Developing Applications with WebLogic Server*.

Note: This mechanism is only for overriding resources and does not override classes.

These are application-specific files and the contents are opaque to WLS, so the entire file contents will be overridden when an override file is supplied.

How It Works

The files placed in the `/AppFileOverrides` subdirectory are staged and distributed along with the rest of the plan directory contents and are available on all of the targets. Applications are then able to load these files as resources using the current `ClassLoader` (for example, using the `ClassLoader.getResourceAsStream` method.) This either finds the overridden files or the files packaged in the application, depending on the configuration and whether overridden files are supplied.

For web applications, application file overrides only apply to the classpath related resources (which are in `WEB-INF/classes` and `WEB-INF/lib`), and *do not apply* to the resource path for the web application. Therefore, overrides are seen by web applications using the `classloader.getResourceAsStream()` method to lookup resources, but overrides do not affect web application calls to the `ServletContext.getResourceAsStream()` method.

In order to use this feature, you must:

- Specify a plan for the deployment (see [Creating a New Deployment Plan to Configure an Application](#)).
- Specify the `config-root` within in the plan.
- Provide a `config-root/AppFileOverrides` subdirectory.

Directory Structure

The contents of the `/AppFileOverrides` subdirectory use the existing plan directory structure and directory naming conventions that already exist for descriptor overrides. For more information on directory naming conventions, see [“Packaging Files for Deployment” on page 3-2](#).

Enabling application file overrides causes a directory `ClassLoader` to be added to the application and module level `ClassLoaders`, which point to the appropriate root directories within the `/AppFileOverrides` subdirectory (which is in the plan directory). The `ClassLoader` inserted into the front of the application’s `ClassLoader` is given a structure of

```
AppDeploymentMBean.getLocalPlanDir + separator + "/AppFileOverrides". The
ClassLoader inserted into the front of the module’s ClassLoaders is given a structure of
AppDeploymentMBean.getLocalPlanDir + separator + "/AppFileOverrides" +
separator + moduleURI.
```

For Example:

Table 4-1 Directory Structure for Generic File Overrides

Directory	Description
<code>install-root/plan/AppFileOverrides</code>	Directory put in front of the main application classloader’s classpath
<code>install-root/plan/AppFileOverrides/WebApp1.war/...</code>	Directory put in front of the <code>WebApp1.war</code> classloader’s classpath
<code>install-root/plan/AppFileOverrides/WebApp2.war/...</code>	Directory put in front of the <code>WebApp2.war</code> classloader’s classpath

Application Usage

It is important to note that the application controls the file contents and format and controls whether the contents of the files are accessed by the application code. As a best practice, generic file loading overrides should be used by application code that has environment-specific properties files, and which is loading those properties files as resources using the application’s classloader. For example, the application code may do the following:

```

Properties myAppProps = new Properties();
InputStream iostream =
Thread.currentThread().getContextClassLoader().getResourceAsStream("myCfg/
myApp.properties");
myAppProps.load(iostream);

```

Additional Configuration Tasks

See the following sections for information about additional deployment configuration tasks:

- [“Deploying an Application with a Deployment Plan” on page 6-4](#) describes how to deploy an application with a valid deployment plan using the `weblogic.Deployer` tool. See [“weblogic.Deployer Command-Line Reference” on page A-1](#).
- [“Updating the Deployment Configuration for an Application” on page 8-22](#) describes how to update the deployment configuration for a currently-deployed application.
- [“Exporting an Application for Deployment to New Environments” on page 5-1](#) explains how developers can create portable deployment plans using the `weblogic.PlanGenerator` tool.
- [“weblogic.PlanGenerator Command Line Reference” on page B-1](#) provides a complete reference to the `weblogic.PlanGenerator` tool.

Best Practices for Managing Application Configuration

- Always manage multiple deployment configurations using deployment plans, rather than multiple versions of the WebLogic Server deployment descriptor files.
- Always store each existing deployment plan for an application in its own `plan` subdirectory of an application root directory.
- If your organization requires standardized, repeatable deployments to several environments, use the [“Application with Single Deployment Plan” on page 4-6](#) workflow to maintain a single deployment plan in your source control system.
- If you make extensive changes to an application’s deployment configuration using the Administration Console, back up or safely store the updated deployment plan for future use. Oracle recommends storing the entire application root directory in a source control system, so that you can maintain configuration information for multiple environments and multiple versions of an application.

Configuring Applications for Production Deployment

Exporting an Application for Deployment to New Environments

The following sections describe how to export an application's WebLogic Server deployment configuration to a custom deployment plan, which helps administrators easily deploy the application into non-development environments:

- [“Overview of the Export Process” on page 5-1](#)
- [“Understanding Deployment Property Classifications” on page 5-3](#)
- [“Steps for Exporting an Application's Deployment Configuration” on page 5-4](#)
- [“Staging Application Files for Export” on page 5-4](#)
- [“Generating a Template Deployment Plan using weblogic.PlanGenerator” on page 5-5](#)
- [“Customizing the Deployment Plan Using the Administration Console” on page 5-6](#)
- [“Manually Customizing the Deployment Plan” on page 5-7](#)
- [“Validating the Exported Deployment Configuration” on page 5-8](#)
- [“Best Practices for Exporting a Deployment Configuration” on page 5-9](#)

Overview of the Export Process

Exporting an application's deployment configuration is the process of creating a custom deployment plan for deploying the application into new WebLogic Server environments. When the process is complete, the application deployment files and the custom deployment plan are

distributed to deployers (for example, testing, staging, or production administrators) who then use the deployment plan as a blueprint for configuring the application for their environment.

An administrator can install the application and the custom deployment plan using the Administration Console, which validates the deployment plan and allows the administrator to update configuration properties needed for a specific deployment.

See the [“Understanding Deployment Plan Contents” on page 4-13](#) for more information about deployment plans.

Goals for Exporting a Deployment Configuration

The primary goals in exporting a deployment configuration are:

- **To expose the external resources requirements of the application as null variables in a deployment plan.** Any external resources required by the application are subject to change when the application is deployed to a different environment. For example, the JNDI names of datasources used in your development environment may be different from those used in testing or production. Exposing those JNDI names as variables makes it easy for deployers to use available resources or create required resources when deploying the application. Using empty (null) variables forces the deployer to fill in a valid resource name before the application can be deployed.
- **To expose additional configurable properties, such as tuning parameters, as variables in a deployment plan.** Certain tuning parameters that are acceptable in a development environment may be unacceptable in a production environment. For example, it may suffice to accept default or minimal values for EJB caching on a development machine, whereas a production cluster would need higher levels of caching to maintain acceptable performance. Exporting selected tunables as deployment plan variables helps an administrator focus on important tuning parameters when deploying the application. The Administration Console highlights tuning parameters exposed as variables in a deployment plan, but does not require a deployer to modify them before deployment.

Tools for Exporting a Deployment Configuration

WebLogic Server provides the following tools to help you 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 (see [“Understanding](#)

[Deployment Property Classifications” on page 5-3](#)). You typically need to manually modify the deployment plan created by `weblogic.PlanGenerator`, either manually or using the Administration Console, to delete extraneous variable definitions or add variables for individual properties.

- 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.

Understanding Deployment Property Classifications

Each WebLogic Server deployment descriptor property (for all J2EE module descriptors as well as JDBC, JMS, and WLDF application modules) can be classified into one of the following categories:

- *Non-configurable* properties cannot be changed by an administrator during a deployment configuration session. Non-configurable properties are used to describe application behavior that is fundamental to the basic operation of the application. For example, the `ejb-name` property is categorized as non-configurable, because changing its value also requires changing the EJB application code.
- *Dependency* properties resolve resource dependencies defined in the J2EE deployment descriptors. For example, if the J2EE descriptor for an EJB defines a datasource name that is used within the EJB code, the WebLogic Server descriptor uses a dependency property to bind the datasource name to an actual datasource configured in the target WebLogic Server domain.
- *Declaration* properties declare a resource that other applications can use. For example, the JNDI name of an EJB declares the EJB name that other applications or modules would use to access the EJB.
- *Configurable* properties are the remaining properties not classified as dependency or declaration properties. Generally configurable properties enable or configure WebLogic Server-specific features and tuning parameters for the deployed application. For example, the WebLogic Server descriptor for an EJB might define the number of EJBs that WebLogic Server caches in memory.

Use these categories during the configuration export process to select properties to expose as variables in the deployment plan. For example, a developer can generate a new deployment plan containing variable definitions for all properties tagged as “dependencies” in an application’s WebLogic Server deployment descriptors. The variables can then be easily changed by an administrator deploying the application to an environment having different resource names.

All changeable descriptor properties (dependency, declaration, and configurable properties) are further classified as either *dynamic* or *non-dynamic* properties. Dynamic properties can be changed in a deployed application without requiring you to redeploy for the changes to take effect. Non-dynamic properties can be changed but require redeployment for the changes to take effect. The Administration Console identifies non-dynamic properties as a reminder for when redeployment is necessary.

Steps for Exporting an Application’s Deployment Configuration

Exporting an application’s deployment configuration typically involves the following procedures:

1. [Staging Application Files for Export](#)
2. [Generating a Template Deployment Plan using `weblogic.PlanGenerator`](#)
3. [Customizing the Deployment Plan Using the Administration Console](#)
4. [Manually Customizing the Deployment Plan](#)
5. [Validating the Exported Deployment Configuration](#)

The sections that follow describe each procedure in detail.

Staging Application Files for Export

Oracle recommends placing application files into an application installation directory before exporting the deployment configuration. When using an installation directory, generated configuration files, such as the deployment plan, are automatically copied to the `\plan` subdirectory during export.

To create an application installation directory:

1. Create a top-level installation directory for your application:

```
mkdir c:\exportapps\myApplication
```

2. Create `\app` and `\plan` subdirectories:

```
mkdir c:\exportapps\myApplication\app
```

```
mkdir c:\exportapps\myApplication\plan
```

3. Copy the complete application to be exported into the `\app` subdirectory. The application can be either in archive or exploded archive form:

```
cp -r c:\dev\myApplication c:\exportapps\myApplication\app
```

The `\app` directory must include the full application distribution, and can include the WebLogic Server descriptor files that you use for deployment to your development environment.

If you choose not to use an installation directory when exporting an application, Oracle recommends using the `-plan` option to `weblogic.PlanGenerator` to specify the location and filename of the generated plan. By default, `weblogic.PlanGenerator` stores generated files in the `TEMP/weblogic-install/application_name/config` directory, where `TEMP` is the temporary directory for your environment. For Windows platforms, this means generated configuration files are stored in `C:\Documents and Settings\username\Local Settings\Temp\weblogic\install\myApplication.ear\config`. Use the `-plan` option to place generated files in a known location.

Generating a Template Deployment Plan using `weblogic.PlanGenerator`

The `weblogic.PlanGenerator` tool provides a quick and easy way to generate a template deployment plan with null variables for an entire category (such as declaration or configurable properties) of deployment descriptors. Oracle recommends using `weblogic.PlanGenerator` to generate a new deployment plan with null variables for all of an application's dependencies. This ensures that all global resources required for an application can be easily configured by administrators who must deploy the application in a new environment.

When using an application staged in an installation root directory, the basic syntax for using `weblogic.PlanGenerator` is:

```
java weblogic.PlanGenerator -root install_root category
```

where:

install_root specifies the fully qualified name of the root directory for the application and plan.

category specifies the category of WebLogic Server deployment descriptors for which you want to create variables. (See [“Understanding Deployment Property Classifications”](#) on page 5-3 for a description of each category.) For the purposes of generating a template deployment plan, you should usually use only the `-dependencies` option, which is the default option, as this limits variables to external resources required by the application.

Note: The `-dependencies` option creates null variables for every possible dynamically-configurable deployment property, which can result in a large number of variable definitions that may not be required for your application. The `-declarations` option is generally not required, because declaration properties are typically associated with the basic functioning of the application and should not be changed before deployment to a new environment.

For example:

```
java weblogic.PlanGenerator -root c:\exportapps\myApplication
-dependencies
```

```
java weblogic.PlanGenerator -root c:\exportapps\myApplication
```

With the above commands, which are synonymous because `-dependencies` is the default option so you are not required to specify it in your `weblogic.PlanGenerator` command, `weblogic.PlanGenerator` inspects all J2EE deployment descriptors in the selected application, and creates a deployment plan with null variables for all relevant WebLogic Server deployment properties that configure external resources for the application. Using this template deployment plan, an administrator using the Administration Console is directed to assign valid resource names and tuning properties for each null variable before the application can be deployed.

Customizing the Deployment Plan Using the Administration Console

The template deployment plan generated in [“Generating a Template Deployment Plan using `weblogic.PlanGenerator`”](#) on page 5-5 contains only those deployment properties that resolve external dependencies for the application. A developer generally customizes the template plan to add one or more WebLogic Server tuning properties for the application. The Administration Console enables you to easily add deployment plan variables for individual deployment descriptor properties as needed. To customize a deployment plan using the Administration Console:

1. [“Install the Exported Application and Template Deployment Plan”](#) on page 5-7

2. [Add Variables for Selected Tuning Properties](#)
3. [Retrieve the Customized Deployment Plan](#)

Install the Exported Application and Template Deployment Plan

To modify a deployment configuration using the Administration Console, you must first install the application and existing deployment plan as described in “[Steps for Creating an Application Installation Directory](#)” on page 3-7.

Add Variables for Selected Tuning Properties

After installing the exported application, follow the steps in [Update a deployment plan](#) in *Administration Console Online Help* to add new tuning properties to the deployment plan.

Retrieve the Customized Deployment Plan

When you modify an application’s deployment configuration using the Administration Console, your changes to deployment properties are stored in a WebLogic Server deployment plan and/or in generated WebLogic Server deployment descriptor files. If you modify any deployment properties defined as variables in the application’s deployment plan, your changes are written back to a new version of the plan file. If the application that was installed from an installation directory, the Administration Console stores the generated configuration files in the `plan` subdirectory by default.

Manually Customizing the Deployment Plan

In some cases you may need to edit a custom deployment plan manually, using a text editor. This may be necessary for the following reasons:

- You want to remove an existing deployment plan variable.
- You want to assign a null value to a generated variable in the plan.

Note: You cannot use the Administration Console to remove variable definitions from the deployment plan or assign a null value for a deployment property.

See <http://www.bea.com/ns/weblogic/deployment-plan/1.0/deployment-plan.xsd> and “[Understanding Deployment Plan Contents](#)” on page 4-13 before manually editing deployment plan entries.

Removing Variables from a Deployment Plan

The `variable-definition` stanza in a deployment plan defines the names and values of variables used for overriding WebLogic Server deployment descriptor properties. The `module-override` element may contain one or more `variable-assignment` elements that define where a variable is applied to a given deployment descriptor. To remove a variable from a deployment plan, use a text editor to delete:

- The `variable` definition from the `variable-definition` stanza
- All `variable-assignment` elements that reference the deleted variable.

Assigning Null Variables to Require Administrator Input

To assign a null value to an existing variable definition, simply change any text value that is present in the `value` subelement in the `variable` element to `<value xsi:nil="true"></value>` where the `xsi` namespace is defined as: `xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"`. For example, change:

```
...  
<variable-definition>  
  <variable>  
    <name>SessionDescriptor_InvalidationIntervalSecs_11029744771850</name>  
    <value>80</value>  
  </variable>  
</variable-definition>  
...
```

to:

```
...  
<variable-definition>  
  <variable>  
    <name>SessionDescriptor_InvalidationIntervalSecs_11029744771850</name>  
    <value xsi:nil="true"></value>  
  </variable>  
</variable-definition>  
...
```

Validating the Exported Deployment Configuration

The Administration Console automatically validates the deployment configuration for a newly-installed application or module. To validate a custom deployment plan that you have created during the export process:

1. Follow the steps under “[Add Variables for Selected Tuning Properties](#)” on page 5-7 to install the application or module with the final version of the custom deployment plan. The Administration Console automatically uses a deployment plan named `plan.xml` in the `plan` subdirectory of an installation directory, if one is available.
2. On the Summary of Deployments page, select the name of the application or module that you installed.
3. Select the Deployment Plan > Dependencies tab.
4. Verify that the dependencies configured for the deployed module are valid for the selected target servers.

Best Practices for Exporting a Deployment Configuration

Keep in mind these best practices when exporting an application’s deployment configuration:

- The primary goal for exporting an application is to create null variables for all of an application’s external resource dependencies. This ensures that deployers have the ability to assign resource names based on resources available in their target environment.
- Use `weblogic.PlanGenerator` only for exporting resource dependencies. Using `weblogic.PlanGenerator` to export other categories of deployment descriptor properties generally results in too many variables in the deployment plan.
- Use the Administration Console to add individual tuning property values to the deployment plan, or to validate a custom deployment plan.
- Neither the Administration Console nor `weblogic.PlanGenerator` allow you to remove variables from a plan or set null values for variables. Use a text editor when necessary to complete these tasks.

Exporting an Application for Deployment to New Environments

Deploying Applications and Modules with `weblogic.deployer`

The following sections describe how administrators and developers perform interactive, command-line based deployment tasks using the `weblogic.Deployer` utility:

- [“Overview of Common Deployment Scenarios” on page 6-2](#)
- [“Deploying to a Single-Server Domain” on page 6-4](#)
- [“Deploying an Application with a Deployment Plan” on page 6-4](#)
- [“Deploying an Application That Looks Up System Resources from JNDI During `preStart`” on page 6-5](#)
- [“Uploading Deployment Files from a Remote Client” on page 6-3](#)
- [“Targeting Deployments to Servers, Clusters, and Virtual Hosts” on page 6-6](#)
- [“Using Module-Level Targeting for Deploying an Enterprise Application” on page 6-9](#)
- [“Deploying JDBC, JMS, and WLDF Application Modules” on page 6-11](#)
- [“Controlling Deployment File Copying with Staging Modes” on page 6-14](#)
- [“Distributing Applications to a Production Environment” on page 6-21](#)
- [“Deploying Shared Java EE Libraries and Dependent Applications” on page 6-22](#)
- [“Best Practices for Deploying Applications” on page 6-26](#)

Overview of Common Deployment Scenarios

The sections that follow organize common deployment tasks into several general categories, starting with the most basic tasks and progressing to more advanced tasks. Before you deploy an application, perform the appropriate tasks described in [“Configuring Applications for Production Deployment”](#) on page 4-1.

- [“Uploading Deployment Files from a Remote Client”](#) on page 6-3—Explains how to upload an application or module to the Administration Server from a remote client.
- [“Deploying to a Single-Server Domain”](#) on page 6-4—Explains the basics of deploying an application or module to a single-server WebLogic domain.
- [“Deploying an Application with a Deployment Plan”](#) on page 6-4—Explains how to deploy an application with a deployment plan that fully configures the application.
- [“Deploying an Application That Looks Up System Resources from JNDI During `preStart`”](#) on page 6-5—Explains how to deploy an application that looks up system resources via JNDI.
- [“Targeting Deployments to Servers, Clusters, and Virtual Hosts”](#) on page 6-6—Describes all WebLogic Server target types, and explains how to identify targets during deployment.
- [“Using Module-Level Targeting for Deploying an Enterprise Application”](#) on page 6-9—Describes how to target individual modules in an Enterprise Application to different WebLogic Server targets.
- [“Deploying JDBC, JMS, and WLDF Application Modules”](#) on page 6-11—Describes how to deploy both stand-alone and application-scoped JDBC, JMS, and WLDF modules in a WebLogic Server domain.
- [“Controlling Deployment File Copying with Staging Modes”](#) on page 6-14—Describes how to use non-default staging modes to control WebLogic Server file-copying behavior.
- [“Distributing Applications to a Production Environment”](#) on page 6-21—Describes how to deploy and test a new application directly in a production environment *without* making the application available for processing client requests.
- [“Deploying Shared Java EE Libraries and Dependent Applications”](#) on page 6-22—Describes how to deploy Java EE Libraries and optional packages to a Weblogic Server environment, and how to deploy and manage applications and modules that reference these shared resources.

- [“Best Practices for Deploying Applications” on page 6-26](#)—Summarizes key deployment practices.

Uploading Deployment Files from a Remote Client

In order to deploy an application or module to a domain, the deployment file(s) must be accessible to the domain’s Administration Server. If the files do not reside on the Administration Server machine or are not available to the Administration Server machine via a network mounted directory, use the `-upload` option to upload the files before deploying them:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -deploy -upload c:\localfiles\myapp.ear
```

To upload an exploded archive directory, specify the directory name instead of an archive filename (for example `c:\localfiles\myappEar`).

When you upload files to the Administration Server machine, the archive file is automatically placed in the server’s upload directory. You can configure the path of this directory using the instructions in [“Changing the Default Staging Behavior for a Server” on page 6-21](#).

Upload Behavior When Updating an Application or Plan

When mixing local and remote (upload) deployment operations, it is very important to use a process that tracks how an application is deployed. Consider the following series of events:

1. An administrator or developer deploys an app (`myapp.ear`) with a deployment plan (`productionEnvPlan.xml`) to an administration server having a managed server using:

```
java weblogic.Deployer -adminurl http://test:7001 -user weblogic
  -password weblogic -deploy c:\localfiles\myapp.ear
  -plan c:\localfiles\productionEnvPlan.xml
```

2. If the administrator or developer uses WebLogic Portal on the managed server and saves some configuration changes to the `myapp.ear` application. When the application is saved, WebLogic Portal updates the application using the following:

```
java weblogic.Deployer -adminurl http://test:7001 -user weblogic
  -password weblogic -update -name myapp.ear -upload
  -plan c:\localfiles\nuPlan.xml
```

3. The application and deployment plan file from the remote machine are uploaded to the administration server’s upload directory. The deployment is updated and the configuration

uses the deployment plan at

```
c:\domain\servers\adminServerName\upload\plan\nuPlan.xml.
```

At this point, an administrator or developer expecting the behavior derived from the application (`c:\localfiles\myapp.ear`) and deployment plan (`c:\localfiles\productionEnvPlan.xml`) can easily become confused. If the administrator or developer:

- Reviews the application and deployment plan files, the files have not changed. New behaviors appear to be unexplained.
- Makes changes to the `c:\localfiles\myapp.ear` application file or `c:\localfiles\productionEnvPlan.xml` deployment plan file `xml`, the deployment configuration changes in the upload directory are at risk of being lost.

Deploying to a Single-Server Domain

A single-server WebLogic Server domain, consisting only of an Administration Server, represents the simplest scenario in which to deploy an application or module. If you are deploying files that reside on the same machine as the domain, use the `-deploy` command and identify the file location, with connection arguments for the Administration Server. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic  
-password weblogic -deploy c:\localfiles\myapp.ear
```

In the above command, WebLogic Server creates a default deployment name of `myapp`, as described in [“Understanding Default Deployment Names” on page 3-4](#), because `myapp` is the name of the deployment file without the extension. If you want to specify a non-default deployment name, use the `-name` option, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001  
-user weblogic -password weblogic -deploy  
-name myTestApplication c:\localfiles\myapp.ear
```

To deploy an application or module to multiple targets, see [“Targeting Deployments to Servers, Clusters, and Virtual Hosts” on page 6-6](#)

Deploying an Application with a Deployment Plan

When you use `weblogic.Deployer` to deploy an application, the deployment plan and WebLogic Server deployment descriptors must define a valid configuration for the target

environment, or the deployment fails. This means you cannot use `weblogic.Deployer` with a deployment plan that defines null variables for an application's required resource bindings.

To deploy an application and deployment plan using `weblogic.Deployer`, include the `-plan` option with the **-deploy** command, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -deploy -name myTestDeployment
  -source /myDeployments/myApplication.ear
  -targets myCluster -stage
  -plan /myDeployments/myAppPlan.xml
```

If you are deploying from an application root directory and the deployment plan is located in the `/plan` subdirectory, you still need to identify both the actual deployment source files the plan to use for deployment, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -deploy -name myTestDeployment
  -source
/myDeployments/installedApps/myApplication/app/myApplication.ear
  -targets myCluster -stage
  -plan /myDeployments/installedApps/myApplication/plan/plan.xml
```

When you deploy or distribute an application with a deployment plan, the deployment plan and any generated deployment descriptors are copied to the staging directories of target servers along with the application source files.

Deploying an Application That Looks Up System Resources from JNDI During preStart

Because `preStart` application lifecycle listeners are invoked during the prepare phase of an application—which occurs when an edit session is activated—JNDI lookup of system resources fails if you create a new system resource during the same edit session in which you deploy an application that uses the system resource. For example, if you acquire an edit lock from the Administration Console, then create a new JDBC system resource and deploy an application that uses the new JDBC system resource during the same edit session (in other words, before activating changes in the Administration Console), JNDI lookup of the JDBC system resource will fail because the `preStart` application lifecycle listeners were invoked before the system resource was activated. To look up system resources from JNDI in your `preStart` lifecycle listeners, Oracle recommends that you create the system resource during startup or during a

separate edit session, before you deploy any applications that use the system resource. For example:

1. Acquire an edit lock by clicking **Lock & Edit** in the Administration Console.
2. Create the desired system resource. For information, see [Create JMS system modules](#) in *Administration Console Online Help*.
3. Complete the edit session by clicking **Activate Changes** in the Administration Console.
4. Acquire another edit lock by clicking **Lock & Edit** in the Administration Console.
5. Deploy the application that uses the system resource you created in step 2. For information on deploying applications from the Administration Console, see [Deploy applications and modules](#) in *Administration Console Online Help*.

See [Overview of Java EE Libraries and Optional Packages](#) in *Developing Applications with WebLogic Server* for more information.

Targeting Deployments to Servers, Clusters, and Virtual Hosts

In most production environments, you typically deploy applications to one or more Managed Servers configured in a Weblogic Server domain. In some cases, the servers may be included as part of a WebLogic Server cluster, or a virtual host may be used for directing Web application requests. The term *deployment target* refers to any server or collection of servers to which you can deploy an application or module.

Understanding Deployment Targets

Deployment targets are the servers or groups of servers to which you deploy an application or stand-alone module. During the deployment process, you select the list of targets from the available targets configured in your domain. You can also change the target list after you have deployed a module.

The following table describes all valid deployment targets for WebLogic Server, and lists the types of modules that you can deploy to each target.

Table 6-1 WebLogic Server Deployment Targets

Target Type	Description	Valid Deployments
WebLogic Server Instance	A WebLogic Server instance, such as an Administration Server in a single-server domain, or a Managed Server.	Java EE Applications Java EE modules JMS, JDBC, or WLDF modules Java EE Libraries
Cluster	A configured cluster of multiple WebLogic Server instances	Java EE Applications Java EE modules JMS, JDBC, or WLDF modules Java EE Libraries
Virtual Host	A configured host name that routes requests for a particular DNS name to a WebLogic Server instance or cluster. See Configuring Virtual Hosting in Designing and Configuring WebLogic Server Environments for more information.	Web applications
JMS Server	A JMS Server configured in a Weblogic Server domain	A JMS queue or topic defined within a JMS module*

*When deployed as a stand-alone application module, a JMS, JDBC, or WLDF resource appears as a Java EE deployment in the Administration Console.

A stand-alone JMS application module can be targeted to server, cluster, or virtual host targets; queues and topics defined within a JMS module can be further targeted to a configured JMS server. For information on sub-module targeting, see “[Using Sub-Module Targeting with JMS Application Modules](#)” on page 6-12.

Deploying to One or More Targets

To deploy to a single WebLogic Server target, list the configured target name after the `-targets` option to `weblogic.Deployer`. For example, to deploy a Web application to a configured virtual host named `companyHost`, use the command:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -deploy -targets companyHost
c:\localfiles\myWebApp.ear
```

Specify multiple targets using a comma-separated list, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -deploy -targets ManagedServer-1,ManagedServer-2
    c:\localfiles\myapp.ear
```

Deploying to a Cluster Target

If you specify a cluster target (using `-targets mycluster`), WebLogic Server targets all server instances in the cluster by default. This corresponds to homogenous module deployment, which is recommended in most clusters. If you want to deploy a module only to a single server in the cluster (that is, “pin” a module to a server), specify the individual server instance name, rather than the cluster, as the target. This type of deployment is less common, and should be used only in special circumstances where pinned services are required. See [Understanding Cluster Configuration and Application Deployment](#) in *Using WebLogic Server Clusters* for more information.

Note: Pinning a deployment to a subset of server instances in a cluster (rather than to a single server on the cluster) is not recommended and will generate a warning message.

When you deploy an application to a cluster target, WebLogic Server ensures that the deployment successfully deploys on all available members of the cluster. If even a single, available WebLogic Server instance in the cluster cannot deploy the application, the entire deployment fails and no servers in the cluster start the application. This helps to maintain homogeneous deployments to the cluster, because deployment operations succeed or fail as a logical unit.

If a clustered server is unreachable at the time of deployment (for example, because of a network failure between the Administration Server and a Managed Server, or because a cluster member is shut down) that server does not receive the deployment request until the network connection is restored. This default behavior ensures that most deployment operations succeed, even when servers are taken offline.

Enforcing Consistent Deployment to All Configured Cluster Members

The default cluster deployment behavior ensures homogeneous deployment for all clustered server instances that can be reached at the time of deployment. However, if the Administration Server cannot reach one or more clustered servers due to a network outage, those servers do not receive the deployment request until the network connection is restored. For redeployment operations, this can lead to a situation where unreachable servers use an older version of the

deployed application, while reachable servers use the newer version. When the network connection is restored, previously-disconnected servers may abruptly update the application as they receive the delayed redeployment request.

It is possible to change WebLogic Server's default deployment behavior for clusters by setting the `ClusterConstraintsEnabled` option when starting the WebLogic Server domain. The `ClusterConstraintsEnabled` option enforces strict deployment for all servers configured in a cluster. A deployment to a cluster succeeds only if all members of the cluster are reachable and all can deploy the specified files.

WARNING: Do not use the `ClusterConstraintsEnabled` option unless you have an extremely reliable network configuration, and you can guarantee that all cluster members are always available to receive deployment and redeployment requests. With `ClusterConstraintsEnabled`, WebLogic Server will fail all deployment operations to a cluster if any clustered server is unavailable, even if a single server has been shut down for maintenance.

To set the `ClusterConstraintsEnabled` for the domain when you start the Administration Server, include the appropriate startup argument:

- `-DClusterConstraintsEnabled=true` enforces strict cluster deployment for servers in a domain.
- `-DClusterConstraintsEnabled=false` ensures that all available cluster members deploy the application or module. Unavailable servers do not prevent successful deployment to the available clustered instances. This corresponds to the default WebLogic Server deployment behavior.

Using Module-Level Targeting for Deploying an Enterprise Application

An Enterprise Application (EAR file) differs from other deployment units because an EAR can contain other module types (WAR and JAR archives). When you deploy an Enterprise Application using the Administration Console, you can target all of the archive's modules together as a single deployment unit, or target individual modules to different servers, clusters, or virtual hosts.

You can also use module-level targeting to deploy only a subset of the modules available in an EAR. This can simplify packaging and distribution of applications by packaging multiple modules in a single, distributable EAR, but targeting only the modules you need to each domain.

Module-Targeting Syntax

To target individual modules in an Enterprise Application, use the *module_name@target_name* syntax. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name myEnterpriseApp
    -targets module1@myserver1,module2@myserver2,module3@myserver3
    -stage -deploy c:\localfiles\myEnterpriseApp.ear
```

Targeting Web Application Modules

To target Web application modules that are part of an .ear file, you can use the Web application's *context-root* name as the module name or specify the *web-uri*.

For example, if the *application.xml* file for a file, *myEnterpriseApp.ear*, defines:

```
<module>
  <web>
    <web-uri>myweb.war</web-uri>
    <context-root>/welcome</context-root>
  </web>
</module>
```

You can deploy only the Web application module by using the *context-root* name:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mywebapplication -targets /welcome@myserver1
    -stage -deploy c:\localfiles\myEnterpriseApp.ear
```

You can deploy only the Web application module by using the *web-uri*:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mywebapplication -targets myweb.war@myserver1
    -stage -deploy c:\localfiles\myEnterpriseApp.ear
```

To deploy a Web application as a default Web application, set the value of the *context-root* element to *"/*. For example, if the *application.xml* file for a file, *myEnterpriseApp.ear*, defines:

```
<module>
  <web>
    <web-uri>myweb.war</web-uri>
    <context-root>/</context-root>
```

```
</web>
</module>
```

You can deploy only the Web application module by using the `context-root` name:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name mywebapplication -targets /@myserver1
  -stage -deploy c:\localfiles\myEnterpriseApp.ear
```

Deploying JDBC, JMS, and WLDF Application Modules

Stand-alone JDBC, JMS, and WLDF application modules can be deployed similar to stand-alone Java EE modules. For a stand-alone JDBC, JMS, or WLDF application module, the target list determines the WebLogic Server domain in which the module is available. JNDI names specified within an application module are bound as global names and available to clients. For example, if you deploy a stand-alone JDBC application module to a single-server target, then applications that require resources defined in the JDBC module can only be deployed to the same server instance. You can target application modules to multiple servers, or to WebLogic Server clusters to make the resources available on additional servers.

If you require JDBC, JMS, or WLDF resources to be available to all servers in a domain, create system modules, rather than deployable application modules. See [Configuring JMS System Resources](#) in *Configuring and Managing WebLogic JMS*, [Configuring WebLogic JDBC Resources](#) in *Configuring and Managing WebLogic JDBC*, and [Configuring Diagnostic System Modules](#) in *Configuring and Using the WebLogic Diagnostic Framework*.

Note: Deploying a JMS module does not necessarily mean it is properly targeted and active, and that the JNDI name is available. The JNDI name is not available until the JMS module is properly targeted. For more information, see [Understanding JMS Resource Configuration](#) in *Configuring and Managing WebLogic JMS*.

Targeting Application-Scoped JMS, JDBC, and WLDF Modules

JMS, JDBC, and WLDF application modules can also be included as part of an Enterprise Application, as an application-scoped resource module. Application-scoped resource modules can be targeted independently of other EAR modules during deployment, if necessary, by using module-level targeting syntax. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name myEnterpriseApp
  -targets myWebApp@myCluster,myJDBCModule@myserver1,myEJBModule@myserver
```

1

```
-stage -deploy c:\localfiles\myEnterpriseApp.ear
```

Using Sub-Module Targeting with JMS Application Modules

Certain JMS resources defined within a JMS application module can be further targeted to a JMS Server available on the module's target during deployment. These resources are known as submodules. Certain types of submodule require deployment to a JMS Server, such as:

- Queues
- Topics

Other submodules can be targeted to JMS Servers as well as WebLogic Server instances and clusters:

- Connection factories
- Foreign servers
- SAF imported destinations
- Uniform distributed topics
- Uniform distributed queues

To specify submodule targets at deployment or undeployment time, you must use an extended form of the module targeting syntax with the `-submoduletargets` option to `weblogic.Deployer`.

Default Submodule Targeting

When deploying from the Administration Console, WebLogic Server selects default JMS Server targets for submodules in a JMS application module, as described in [Targeting JMS Modules and Subdeployment Resources](#) in *Configuring and Managing WebLogic JMS*.

When deploying submodule using WLST, you can use the parent module's targets as default targets for JMS resources when all the following conditions are met:

- The application contains one or more JMS modules.
- There is exactly one JMS server instance associated with the module target.
- The `-submoduletargets` option is not specified

- The `defaultSubmoduleTargets` option is true.

See [Deployment Commands](#) in *WebLogic Scripting Tool*.

Sub-module Targeting for Stand-alone JMS Modules

For a stand-alone JMS module, the submodule targeting syntax is: `-submoduletargets submodule_name@target_name`. For example, to deploy a stand-alone JMS module on a single server instance, targeting the submodule `myQueue` to a JMS Server named `JMSServer1`, enter the command:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name myJMSModule
  -targets ManagedServer1 -submoduletargets myQueue@JMSServer1
  -deploy c:\localfiles\myJMSModule.xml
```

To undeploy the same JMS module, enter the following command, which, assuming `myJMSModule` contains more than one submodule, will undeploy only the `myQueue` submodule; all other submodules are unaffected.

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name myJMSModule
  -undeploy -submoduletargets myQueue@JMSServer1
```

Sub-module Targeting for Application-scoped JMS Modules

For an application-scoped JMS module in an EAR, use the syntax:

`submodule_name@module_name@target_name` to target a submodule. For example, if the queue in the above example were instead packaged as part of an enterprise application, you would use the command:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name myEnterpriseApp
  -targets ManagedServer1 -submoduletargets
myQueue@myJMSModule@JMSServer1
  -deploy c:\localfiles\myEnterpriseApp.ear
```

To undeploy the same JMS module, enter the following command, which, assuming `myEnterpriseApp` contains more than one submodule, will undeploy only the `myQueue` submodule; all other submodules are unaffected.

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name myEnterpriseApp
  -undeploy -submoduletargets myQueue@myJMSModule@JMSServer1
```

WARNING: When you undeploy an application that contains application-scoped resources, the resources are deleted along with the application, which can potentially cause abandoned transactions or lost messages as a result of deleted JMS destinations. For more information, see [Unregister Resource Grace Period](#) in *Programming WebLogic JTA*.

You should only undeploy applications that you are certain you want to completely remove; to temporarily stop client access to applications, use the `-stop` command, described in [weblogic.Deployer Command-Line Reference](#).

For more information on JMS subdeployments, see [Targeting JMS Modules and Subdeployment Resources](#) in *Configuring and Managing WebLogic JMS*.

Controlling Deployment File Copying with Staging Modes

The deployment *staging mode* determines how deployment files are made available to target servers that must deploy an application or stand-alone module. WebLogic Server provides three different options for staging files: stage mode, nostage mode, and external_stage mode.

- [“Staging Mode Descriptions and Best Practices”](#) on page 6-15
- [“Using Nostage Mode Deployment”](#) on page 6-17
- [“Syntax for Stage Mode”](#) on page 6-19 [“Using External_stage Mode Deployment”](#) on page 6-19
- [“Syntax for external_stage Mode”](#) on page 6-20
- [“Changing the Default Staging Behavior for a Server”](#) on page 6-21

Staging Mode Descriptions and Best Practices

The following table describes the behavior and best practices for using the different deployment staging modes.

Table 6-2 Application Deployment Staging Modes

Deployment Staging Mode	Behavior	When to Use
stage	<p>The Administration Server first copies the deployment unit source files to the staging directories of target servers. (The staging directory is named <code>stage</code> by default, and it resides under the target server's root directory.)</p> <p>The target servers then deploy using their local copy of the deployment files.</p>	<ul style="list-style-type: none"> • Deploying small or moderate-sized applications to multiple WebLogic Server instances. • Deploying small or moderate-sized applications to a cluster.

Table 6-2 Application Deployment Staging Modes

Deployment Staging Mode	Behavior	When to Use
nostage	<p>The Administration Server does not copy deployment unit files. Instead, all servers deploy using the same physical copy of the deployment files, which must be directly accessible by the Administration Server and target servers.</p> <p>With nostage deployments of exploded archive directories, WebLogic Server automatically detects changes to a deployment's JSPs or Servlets and refreshes the deployment. (This behavior can be disabled if necessary.)</p>	<ul style="list-style-type: none"> • Deploying to a single-server domain. • Deploying to a cluster on a multi-homed machine. • Deploying very large applications to multiple targets or to a cluster where deployment files are available on a shared directory. • Deploying exploded archive directories that you want to periodically redeploy after changing content. • Deployments that require dynamic update of selected Deployment Descriptors via the Administration Console.

Table 6-2 Application Deployment Staging Modes

Deployment Staging Mode	Behavior	When to Use
external_stage	<p>The Administration Server does not copy deployment files. Instead, the Administrator must ensure that deployment files are distributed to the correct staging directory location before deployment (for example, by manually copying files prior to deployment).</p> <p>With external_stage deployments, the Administration Server requires a copy of the deployment files for validation purposes. Copies of the deployment files that reside in target servers' staging directories are not validated before deployment.</p> <p>You can use the <code>-noversion</code> option to turn off the requirement that deployment files be on the Administration Server, but the <code>-noversion</code> option causes versioning information to be ignored; therefore, you cannot use the <code>-noversion</code> option with versioned applications. For information, see “Common Arguments” on page A-5.</p>	<ul style="list-style-type: none"> • Deployments where you want to manually control the distribution of deployment files to target servers. • Deploying to domains where third-party applications or scripts manage the copying of deployment files to the correct staging directories. • Deployments that do not require dynamic update of selected Deployment Descriptors via the Administration Console (not supported in external_stage mode). • Deployments that do not require partial redeployment of application components.

A server's staging directory is the directory in which the Administration Server copies deployment files for stage mode deployments. It is also the directory in which deployment files must reside before deploying an application using external_stage mode.

Most deployments use either stage or nostage modes, and the WebLogic Server deployment tools use the appropriate default mode when you deploy an application or module. The sections that follow explain how to explicitly set the deployment staging mode when deploying an application or module.

Using Nostage Mode Deployment

In nostage mode, the Administration Server does not copy the archive files from their source location. Instead, each target server must access the archive files from a single source directory for deployment. The staging directory of target servers is ignored for nostage deployments.

For example, if you deploy a Java EE Application to three servers in a cluster, each server must be able to access the same application archive files (from a shared or network-mounted directory) to deploy the application.

Note: The source for the deployment files in nostage mode is the path provided by the user at deployment time (as opposed to stage mode, where the source is the path in each server's staging directory). However, even in nostage mode, WebLogic Server copies out parts of the deployment to temporary directories. This enables users to update entire archived deployments or parts of archived deployments.

In nostage mode, the Web application container automatically detects changes to JSPs and servlets. Nostage also allows you to later update only parts of an application by updating those parts in one file system location and then redeploying.

The Administration Console uses nostage mode as the default when deploying only to the Administration Server (for example, in a single-server domain). `weblogic.Deployer` uses the target server's staging mode, and Administration Servers use nostage mode by default. You can also select nostage mode if you run a cluster of server instances on the same machine, or if you are deploying very large applications to multiple machines that have access to a shared directory. Deploying very large applications in nostage mode saves time during deployment because no files are copied.

Syntax for Nostage Mode

To use nostage mode, specify `-nostage` as an option to `weblogic.Deployer`, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mydeploymentname
    -targets myserver1,myserver2,myserver3 -nostage
    -deploy c:\localfiles\myapp.ear
```

Using Stage Mode Deployment

In stage mode, the Administration Server copies the deployment files from their original location on the Administration Server machine to the staging directories of each target server. For example, if you deploy a Java EE Application to three servers in a cluster using stage mode, the Administration Server copies the deployment files to directories on each of the three server machines. Each server then deploys the Java EE Application using its local copy of the archive files.

When copying files to the staging directory, the Administration Server creates a subdirectory with the same name as the deployment name. So if you deployed using the command:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name mytestear -stage -targets mycluster
  -deploy c:\bea\wlserver_10.3\samples\server\medrecd\dist\physicianEar
```

a new directory, `mytestear`, would be created in the staging directory of each server in `mycluster`. If you do not specify a deployment name, a default deployment name (and staging subdirectory) is used:

- For exploded archive deployments, the deployment name and staging subdirectory are the name of the directory you deployed (*physicianEar* in the example above).
- For archived deployments, the default deployment name is the name of the archive file without the extension. For example, if you deploy `physicianEar.ear`, the deployment name and staging subdirectory are `physicianEar`.

The Administration Console uses stage mode as the default mode when deploying to more than one WebLogic Server instance. `weblogic.Deployer` uses the target server's staging mode as the default, and Managed Servers use stage mode by default.

Stage mode ensures that each server has a local copy of the deployment files on hand, even if a network outage makes the Administration Server unreachable. However, if you are deploying very large applications to multiple servers or to a cluster, the time required to copy files to target servers can be considerable. Consider nostage mode to avoid the overhead of copying large files to multiple servers.

Syntax for Stage Mode

To use stage mode, specify `-stage` as an option to `weblogic.Deployer`, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -name mydeploymentname
  -targets myserver1,myserver2,myserver3 -stage
  -deploy c:\localfiles\myapp.ear
```

Using External_stage Mode Deployment

`External_stage` mode is similar to stage mode, in that target servers deploy using local copies of the deployment files. However, the Administration Server does not automatically copy the deployment files to targeted servers in `external_stage` mode; instead, you must copy the files to the staging directory of each target server before deployment. You can perform the copy manually or use automated scripts.

Within each target server's staging directory, deployment files must be stored in a subdirectory that reflects the deployment name. This can either be the name you type in for the deployment, or the default deployment name (the name of the exploded archive directory, or the name of the archive file without its file extension).

`External_stage` mode is the least common deployment staging mode. It is generally used only in environments that are managed by third-party tools that automate the required copying of files. You may also choose to use `external_stage` mode when you are deploying very large applications to multiple machines and you do not have a shared file system (and cannot use `nostage` mode). Using `external_stage` in this scenario decreases the deployment time because files are not copied during deployment.

You can use the `-noversion` option to turn off the requirement that deployment files be on the Administration Server, but the `-noversion` option causes versioning information to be ignored so you cannot use the `-noversion` option with versioned applications. For information, see [“Common Arguments” on page A-5](#).

Syntax for `external_stage` Mode

To deploy an application in `external_stage` mode:

1. Make sure that the deployment files are accessible to the Administration Server. (See [“Uploading Deployment Files from a Remote Client” on page 6-3](#).)
2. On each target server for the deployment, create a subdirectory in the staging directory that has the same name as the deployment name. For example, if you will specify the name `myEARExternal` for the deployment name, create a `myEARExternal` subdirectory in the staging directories for each target server.
Note: If you do not specify a deployment name at deployment time, WebLogic Server selects a default name. See [“Understanding Default Deployment Names” on page 3-4](#) for more information.
3. If you are deploying a versioned application for use with production redeployment, create a subdirectory beneath the application directory for the version of the application you are deploying. For example, if you are deploying `myEARExternal` at version level `2.0Beta`, the deployment files must reside in a subdirectory of each target server's staging directory named `myEARExternal\2.0Beta`.
4. Copy the deployment files into the staging subdirectories you created in Step 2 or 3 above.
5. Deploy the application or module using the `weblogic.Deployer` utility. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -name weblogic
    -password weblogic -external_stage -name myEARExternal
    -deploy c:\myapps\myear
```

Changing the Default Staging Behavior for a Server

The server staging mode specifies the default deployment mode for a server if none is specified at deployment time. For example, the server staging mode is used if you deploy an application or stand-alone module using `weblogic.Deployer` and you do not specify a staging mode.

Notes: You can only change the server staging mode by using the Administration Console or by directly changing the `ServerMBean` via JMX.

Changing the server staging mode does not affect existing applications. If you want to change the staging mode for an existing application, you must undeploy the application deployment and then redeploy it with the new staging mode.

To set the server staging mode using the Administration Console, follow the steps described in [Set a server staging mode](#) in *Administration Console Online Help*. For detailed information on staging modes, see [Table 6-2, “Application Deployment Staging Modes,”](#) on page 6-15.

Distributing Applications to a Production Environment

Distributing an application prepares it for deployment by copying its deployment files to all target servers and validating it. After you distribute an application, you can start it in Administration mode, which restricts access to the application to a configured Administration channel and allows you to distribute the application to a production environment (or distribute a new version of an application) without opening the application to external client connections.

While in Administration mode, you can connect to an application only via a configured Administration channel. This allows you to perform final (“sanity”) checking of the application directly in the production environment without disrupting clients. After performing final testing, you can either undeploy the application to make further changes, or start the application to make it generally available to clients.

You can use the `-adminmode` option application in Administration mode. For information, see [“Starting a Distributed Application in Administration Mode”](#) on page 6-22.

When deploying across WebLogic domains, the `CrossDomainConnector` role allows you to make inter-domain calls from foreign domains. For a complete listing of all security roles, see [Default Global Roles](#) in *Securing WebLogic Resources Using Roles and Policies*.

Distributing an Application

To distribute an application, use the `weblogic.Deployer` **-distribute** command, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -distribute -name myTestDeployment
    /myDeployments/myApplication/
```

Starting a Distributed Application in Administration Mode

After WebLogic Server distributes the deployment files, you can start the application in Administration mode so that you can access and test it via a configured Administration channel. To configure an administration channel, see [Configuring Network Resources](#) in *Designing and Configuring WebLogic Server Environments*.

To start a distributed application in Administration mode, use the `-start` command with the `-adminmode` option:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -start -adminmode -name myTestDeployment
    /myDeployments/myApplication/
```

Starting a Distributed Application

After performing final testing of a distributed application using a configured Administration channel, you can open the application to new client connections by using the `weblogic.Deployer` `-start` command without the `-adminmode` option:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -start -name myTestDeployment
```

Deploying Shared Java EE Libraries and Dependent Applications

Java EE library support in WebLogic Server provides an easy way to share one or more Java EE modules or JAR files among multiple Enterprise Applications. A Java EE library is a stand-alone Java EE module, multiple Java EE modules packaged in an Enterprise Application (EAR), or a plain JAR file that is registered with the Java EE application container upon deployment. After a Java EE library has been registered, you can deploy Enterprise Applications that reference the library. Each referencing application receives a copy of the shared Java EE library module(s) on deployment, and can use those modules as if they were packaged as part of the application itself.

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.

- [“Understanding Deployment Behavior for Shared Libraries” on page 6-23](#)
- [“Registering Libraries with WebLogic Server” on page 6-24](#)
- [“Specifying Library Versions at Deployment” on page 6-24](#)
- [“Deploying Applications That Reference Libraries” on page 6-25](#)

Understanding Deployment Behavior for Shared Libraries

WebLogic Server supports shared Java EE libraries by merging the shared files with a referencing application when the referencing application is deployed.

You first register a Java EE library or optional package with one or more WebLogic Server instances or clusters by deploying the library and indicating that the deployment is a library (see [“Registering Libraries with WebLogic Server” on page 6-24](#)). Deploying a library or package does not activate the deployment on target servers. Instead, WebLogic Server distributes the deployment files to target servers (if nostage deployment mode is used) and records the location of the deployment files, the deployment name, and any version string information for the library or package.

When an application that references a shared library or package is deployed, WebLogic Server checks the names and version string requirements against the libraries registered with the server. If an exact match for a library or package name is not found, or if the version requirements are not met, the application deployment fails.

If WebLogic Server finds a name and version string match for all of the libraries referenced in the application, the server adds the libraries’ classes to the classpath of the referencing application and merges deployment descriptors from both the application and libraries in memory. The resulting deployed application appears as if the referenced libraries were bundled with the application itself.

The contents of a Java EE library are loaded into classloaders in the same manner as any other Java EE modules in an enterprise application. For example, EJB modules are loaded as part of the referencing application’s classloader, while Web application modules are loaded in classloaders beneath the application classloader. If a shared Java EE library consists of an EAR, any classes stored in the EAR’s `APP-INF/lib` or `APP-INF/classes` subdirectory are also available to the referencing application.

You cannot undeploy any Java EE libraries or optional packages that are referenced by a currently deployed application. If you need to undeploy a shared library or package, you must first undeploy all applications that use the shared files. For regular application maintenance, you should deploy a new version of a shared library or package and redeploy referencing applications to use the newer version of the shared files. See [Editing Manifest Entries for Shared Libraries](#) in *Developing Applications with WebLogic Server* for more information.

Registering Libraries with WebLogic Server

A shared Java EE library is a standard Java EE module or Enterprise Application that is registered with a WebLogic Server container as a library. To register a Java EE library with a WebLogic Server container, perform the following steps:

1. Ensure that the deployment file(s) you are working with represent a valid Java EE library or optional package. See [Creating Shared Java EE Libraries](#) in *Developing Applications with WebLogic Server*.
2. Select the WebLogic Server targets to which you will register the library or package. Shared libraries must be registered to the same WebLogic Server instances on which you plan to deploy referencing applications. (You may consider deploying libraries to all servers in a domain, so that you can later deploy referencing applications as needed.)
3. Register the library or package by deploying the files to your selected target servers, and identifying the deployment as a library or package with the `-library` option. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -deploy -targets myserver1,myserver2
    -library /deployments/myLibraryApplication/
```

Specifying Library Versions at Deployment

As a best practice, your development team should always include version string information for a library or optional package in the manifest file for the deployment. See [Editing Manifest Attributes for Shared Libraries](#) in *Developing Applications with WebLogic Server* for more information about version string requirements and best practices.

If you are deploying a library or package that does not include version string information, you can specify it at the command line using one or both of the following options:

- `libspecver`—Defines a specification version for the library or package.
- `libimplver`—Specifies an implementation version for the library or package.

For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -deploy -targets myserver1,myserver2
    -library -libspevver 700 -libimplversion 7.0.0.1Beta
    /deployments/myLibraryApplication/
```

Notes: If both the manifest file and the `weblogic.Deployer` command line specify version information, and the values do not match, the deployment fails.

If you initially registered a library without providing a specification or implementation version, you must undeploy the library before registering a newer version and specifying version string information.

You can register multiple versions of a library or package, but each version must use the same version strings. For example, you cannot register a library having only a specification version, and then register a new version of the library having both a specification and an implementation version.

Deploying Applications That Reference Libraries

After you deploy the required library and package versions, you can deploy Enterprise applications and modules that reference the libraries. Successfully deploying a referencing application requires that two conditions are met:

- All referenced libraries are registered on the application's target servers.
- Registered libraries meet the version requirements of the referencing application.

There is no special syntax for deploying a referencing application. Deploy the application as you would a standard Enterprise Application or Java EE module.

Note: Referencing applications can be configured with varying levels of version requirements for shared libraries.

An application might be configured to require on of the following:

- The latest version of a shared library or package (the latest registered version), or
- A minimum version of a library or package (or a higher version), or
- An exact version of a library or package.

If you cannot deploy a referencing application because of version requirements, try registering the required version of the conflicting library or package. Or, consult with your development

team to determine whether the version requirements of the application can be relaxed. See [Referencing Libraries in an Enterprise Application](#) in *Developing Applications with WebLogic Server* for more information.

Note: If you have an EAR library that contains Web application components, you cannot deploy multiple applications that reference the library because attempting to do so will result in a context path conflict. This is because `context-root` cannot be set for individual Web applications in EAR libraries; it can only be set for an entire library.

See [Overview of Java EE Libraries and Optional Packages](#) in *Developing Applications with WebLogic Server* for more information.

Best Practices for Deploying Applications

Oracle recommends the following best practices when deploying applications:

- The Administration Server in a multiple-server domain should be used only for administration purposes. Although you can deploy to the Administration Server in a multiple-server domain, this practice is not recommended except during development.
- Package deployment files in an archive format (`.ear`, `.jar`, `.war`, and so forth) when distributing files to different users or environments.
- Check the scenarios described under [“Packaging Files for Deployment”](#) on page 3-2 before deploying. In many cases it is more convenient to deploy applications in exploded format rather than archive format.
- The Administration Console, `weblogic.Deployer` tool, `wldeploy` Ant task, and WLST all provide similar functionality for deploying applications:
 - Use the Administration Console for interactive deployment sessions where you do not know the exact location of deployment files, or the exact names of target servers or deployed applications.
 - Use `weblogic.Deployer` to integrate deployment commands with existing administrative shell scripts or automated batch processes.
 - Use `wldeploy` in conjunction with the split development directory for developing and deploying new applications. `wldeploy` can also be used in place of `weblogic.Deployer` in administration environments that use Ant, rather than shell scripts.
 - Use WLST when you want to create automated scripts that perform deployment tasks.

- Use `wldeploy`, rather than the `autodeploy` directory, to deploy your own applications during development. The `autodeploy` directory is best used for quickly running new applications in a test or temporary environment. For example, if you download a sample application and want to evaluate its functionality, the `autodeploy` directory provides a quick way to deploy the application to a development server.

Deploying Applications and Modules with `weblogic.deployer`

Auto-Deploying Applications in Development Domains

Auto-deployment is a method for quickly deploying an application to a stand-alone server (Administration Server) for evaluation or testing. It is recommended that this method be used only in a single-server development environment.

Note: Oracle recommends that you use the WebLogic Server split development directory and `wldeploy` ant task, rather than auto-deployment, when developing an application. See [Creating a Split Development Directory](#) in *Developing Applications with WebLogic Server*.

If auto-deployment is enabled, when an application is copied into the `\autodeploy` directory of the Administration Server, the Administration Server detects the presence of the new application and deploys it automatically (if the Administration Server is running). If WebLogic Server is not running when you copy the application to the `\autodeploy` directory, the application is deployed the next time the WebLogic Server Administration Server is started. Auto-deployment deploys only to the Administration Server.

Notes: Due to the file locking limitations of Windows, if an application is exploded, all the modules within the application must also be exploded. In other words, you cannot use auto-deployment with an exploded application or module that contains a JAR file.

Auto-deployment is intended for use with a single server target in a development environment. If you use other tools, such as the Administration Console, to add targets to an auto-deployed, exploded application, redeploying the application does not propagate changes to the new target servers.

The following sections provide information on how to use auto-deployment in a development domain:

- “Enabling and Disabling Auto-Deployment” on page 7-2
- “Auto-Deploying, Redeploying, and Undeploying Archived Applications” on page 7-2
- “Auto-Deploying, Redeploying, and Undeploying Exploded Archives” on page 7-3

Enabling and Disabling Auto-Deployment

You can run a WebLogic Server domain in two different modes: development and production. Only development mode allows you use the auto-deployment feature

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). Production mode disables the auto-deployment feature and prevents any applications you place in the `autodeploy` directory after you switch to production mode from being deployed. When you switch from development mode to production mode, any applications that were previously deployed via the `autodeploy` directory remain deployed; if you wish to undeploy or redeploy such applications after you have switched to production mode, you must undeploy or redeploy them manually (for instance, with the `weblogic.Deployer` command and the `-undeploy` or `-redeploy` options, as described in [weblogic.Deployer Command-Line Reference](#)).

By default, a WebLogic Server domain runs in development mode. To specify the mode for a domain, see [Creating and Configuring Domains Using the Configuration Wizard](#).

Auto-Deploying, Redeploying, and Undeploying Archived Applications

To auto-deploy an archived application, copy its archive file to the `/autodeploy` directory. WebLogic Server automatically sets the application’s deployment mode to stage mode.

A deployment unit that was auto-deployed can be dynamically redeployed while the server is running. To dynamically redeploy, copy the new version of the archive file over the existing file in the `/autodeploy` directory.

To undeploy an archived deployment unit that was auto-deployed, delete the application from the `/autodeploy` directory. WebLogic Server stops the application and removes it from the configuration.

Note: If you delete an application from the `/autodeploy` directory when the server is not active, WebLogic Server will not detect that the application was deleted even when the

server is again in an active state. In order to prevent an out-of-sync domain tree, Oracle recommends that you only remove applications from the `/autodeploy` directory when the server is in an active state.

Auto-Deploying, Redeploying, and Undeploying Exploded Archives

To auto-deploy an application in exploded archive format, copy the entire exploded archive directory to the `/autodeploy` directory. WebLogic Server automatically deploys exploded archive applications using the nostage deployment mode.

Notes: Due to Windows file locking limitations, if you deploy an exploded EAR directory that contains archived modules (JAR files), the JAR files are locked during the deployment and you will not be able to remove them. Therefore, if an application you plan to auto-deploy is exploded, all of the modules it contains must also be exploded.

Auto-deployment will fail if you attempt to deploy an application in exploded archive format when the contents of the entire exploded archive directory are not in the `/autodeploy` directory. Because large applications in exploded archive format can take a long time to copy into the `/autodeploy` directory, Oracle recommends that you copy the exploded archive directory into the `/autodeploy` directory while the server is in an inactive state. After the entire exploded archive directory has been copied in the `/autodeploy` directory, you can return the server to an active state and the application will be auto-deployed. Alternatively, Oracle recommends that you deploy large applications with `weblogic.Deployer`, described in [“weblogic.Deployer Command-Line Reference” on page A-1](#).

When an application has been auto-deployed in exploded archive format, the Administration Server periodically looks for a file named `REDEPLOY` in the exploded application directory. If the timestamp on this file changes, the Administration Server redeploys the exploded directory.

To redeploy files in an exploded application directory:

1. When you first deploy the exploded application directory, create an empty file named `REDEPLOY`, and place it in the `WEB-INF` or `META-INF` directory, depending on the application type you are deploying:

An exploded enterprise application has a `META-INF` top-level directory which contains the `application.xml` file.

An exploded Web application has a `WEB-INF` top-level directory which contains the `web.xml` file.

An exploded EJB application has a `META-INF` top-level directory which contains the `ejb-jar.xml` file.

An exploded connector has a `META-INF` top-level directory which contains the `ra.xml` file.

Note: The `REDEPLOY` file works only for an entire deployed application or stand-alone module. If you have deployed an exploded Enterprise Application, the `REDEPLOY` file controls redeployment for the entire application—not for individual modules (for example, a Web application) within the Enterprise Application. If you deploy a Web application by itself as an exploded archive directory, the `REDEPLOY` file controls redeployment for the entire Web application.

2. To redeploy the exploded application, copy the updated files over the existing files in that directory.
3. After copying the new files, modify the `REDEPLOY` file in the exploded directory to alter its timestamp.

When the Administration Server detects the changed timestamp, it redeploys the contents of the exploded directory.

Note: You must `touch` the `REDEPLOY` file (alter its timestamp) any time you wish to trigger redeployment of an auto-deployed application. Even if you modify an application while a server is shut down, you must `touch REDEPLOY` to ensure that changes are applied when the server next starts up.

To undeploy an application that was auto-deployed in exploded format, use the `weblogic.Deployer -undeploy` command, or use the Administration Console to remove the deployment configuration. Then remove the application files from the `/autodeploy` directory.

Note: If you delete application files from the `/autodeploy` directory when the server is not active, WebLogic Server will not detect that the application files were deleted even when the server is again in an active state. In order to prevent an out-of-sync domain tree, Oracle recommends that you only remove application files from the `/autodeploy` directory when the server is in an active state.

Redeploying Applications in a Production Environment

The following sections describe how to use WebLogic Server redeployment to update applications and parts of applications in a production environment:

- [“Overview of Redeployment Strategies” on page 8-1](#)
- [“Understanding When to Use Different Redeployment Strategies” on page 8-4](#)
- [“Using Production Redeployment to Update Applications” on page 8-5](#)
- [“Distributing a New Version of a Production Application” on page 8-13](#)
- [“Using In-Place Redeployment for Applications and Stand-alone Modules” on page 8-16](#)
- [“Using Partial Redeployment for J2EE Module Updates” on page 8-19](#)
- [“Updating Static Files in a Deployed Application” on page 8-21](#)
- [“Updating the Deployment Configuration for an Application” on page 8-22](#)

Overview of Redeployment Strategies

In a production environment, deployed applications frequently require 24x7 availability in order to provide uninterrupted services to customers and internal clients. WebLogic Server provides flexible redeployment strategies to help you update or repair production applications based on their required level of availability.

Production Redeployment

Production redeployment strategy involves 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 version of the application until they complete their work, at which point WebLogic Server automatically retires the older application.

Production redeployment enables you to update and redeploy an application in a production environment without stopping the application or otherwise interrupting the application's availability to clients. Production redeployment saves you the trouble of scheduling application downtime, setting up redundant servers to host new application versions, manually managing client access to multiple application versions, and manually retiring older versions of an application (see [“Redeploying Applications and Modules In-Place” on page 8-18](#) for information about these manual steps).

Production redeployment can be used in combination with the `-distribute` command to prepare a new version of an application for deployment. Then, you can deploy the application in Administration mode, which allows you to perform final sanity testing of a new application version directly in the production environment before making it available to clients. See [“Distributing a New Version of a Production Application” on page 8-13](#).

Production redeployment is supported only for certain J2EE application types, see [“Restrictions on Production Redeployment Updates” on page 8-8](#). For production redeployment procedures and related information, see [“Using Production Redeployment to Update Applications” on page 8-5](#).

In-Place Redeployment

In-place redeployment immediately replaces a running application's deployment files with updated deployment files. In contrast to production redeployment, in-place redeployment of an application or stand-alone J2EE module does not guarantee uninterrupted service to the application's clients. This is because WebLogic Server immediately removes the running classloader for the application and replaces it with a new classloader that loads the updated application class files.

Note: In-place redeployment is the redeployment strategy used in previous versions of WebLogic Server.

WebLogic Server uses the in-place redeployment strategy for J2EE applications that do not specify a version identifier, and for J2EE applications and stand-alone modules that are not supported in [“Production Redeployment” on page 8-2](#). You should ensure that in-place redeployment of applications and stand-alone J2EE modules takes place only during scheduled downtime for an application, or when it is not critical to preserve existing client connections to an application. See [“Using In-Place Redeployment for Applications and Stand-alone Modules” on page 8-16](#) for more information.

WebLogic Server uses the in-place redeployment strategy when performing partial redeployment of a deployed application or module. Partial redeployment can replace either static files in an application, or entire J2EE modules within an Enterprise Application deployment, as described in [“Partial Redeployment of Static Files” on page 8-3](#).

Partial Redeployment of Static Files

WebLogic Server enables you to redeploy selected files in a running application, rather than the entire application at once. This feature is generally used to update static files in a running Web application, such as graphics, static HTML pages, and JSPs. Partial redeployment is available only for applications that are deployed using an exploded archive directory.

Partial redeployment of static files does not affect existing clients of the application. WebLogic Server simply replaces the static files for the deployed application, and the updated files are served to clients when requested. See [“Updating Static Files in a Deployed Application” on page 8-21](#).

Partial Redeployment of J2EE Modules

Partial redeployment also enables you to redeploy a single module or subset of modules in a deployed Enterprise Application. Again, partial deployment is supported only for applications that are deployed using an exploded archive directory.

Note that redeployment for modules in an Enterprise Application uses the in-place redeployment strategy, which does not guarantee uninterrupted client access to the module. For this reason, you should ensure that partial redeployment of J2EE modules in an EAR takes place only during scheduled application downtime, or when it is not critical to preserve client access to the application. See [“Using Partial Redeployment for J2EE Module Updates” on page 8-19](#).

Understanding When to Use Different Redeployment Strategies

The following table summarizes each WebLogic Server redeployment strategy and describes the scenarios in which you would use each strategy.

Table 8-1 Summary of Redeployment Strategies

Redeployment Strategy	Summary	Usage
Production Redeployment	Redeploys a newer version of an application alongside an existing version of the application.	<ul style="list-style-type: none"> Upgrading Web applications and Enterprise Applications that demand uninterrupted client access.
In-Place Redeployment of Applications and Modules	Application classloaders are immediately replaced with newer classloaders to load the updated application class files. WebLogic Server does not guarantee uninterrupted client access during redeployment, and existing clients' state information may be lost.	<ul style="list-style-type: none"> Replacing applications that have been taken off-line for scheduled maintenance. Upgrading applications that do not require uninterrupted client access.
Partial Redeployment of Static Files (In-Place Redeployment)	HTML, JSPs, Graphics Files, or other static files are immediately replaced with updated files.	<ul style="list-style-type: none"> Updating individual Web application files that do not affect application clients.
Partial Redeployment of J2EE Modules (In-Place Redeployment)	Module classloaders are immediately replaced with newer classloaders to load the updated class files. WebLogic Server does not guarantee uninterrupted client access to the module during redeployment, and existing clients' state information may be lost.	<ul style="list-style-type: none"> Replacing a component of an Enterprise Application that has been taken off-line for scheduled maintenance, or that does not require uninterrupted client access.

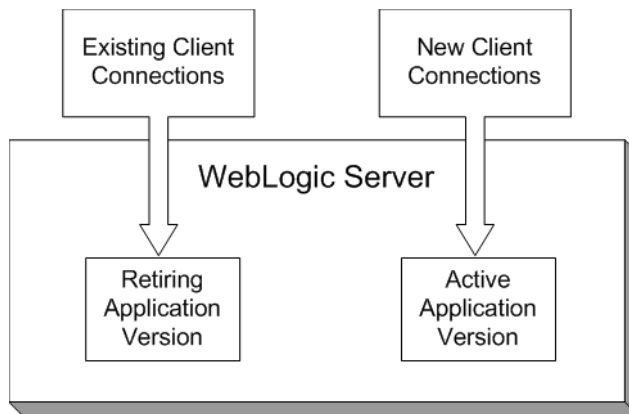
Using Production Redeployment to Update Applications

WebLogic Server enables you to redeploy a new, updated version of a production application without affecting existing clients of the application, and without interrupting the availability of the application to new client requests.

How Production Redeployment Works

WebLogic Server performs production redeployment by deploying a new version of an application alongside an older, running version of the application. While redeployment is taking place, one version of the application “active,” while the other version is “retiring.” The active application version receives all new client connection requests for the application, while the retiring application version processes only those client connections that existed when redeployment took place. WebLogic Server undeploys the retiring application version after all existing clients of the application have finished their work, or when a configured timeout is reached.

Figure 8-1 Production Redeployment



When you redeploy a new version of an application, WebLogic Server treats the newly-deployed application version as the active version, and begins retiring the older version. During the retirement period, WebLogic Server automatically tracks the application’s HTTP sessions and in-progress transactions. WebLogic Server tracks each HTTP session until the session completes or has timed out. In-progress transactions are tracked until the transaction completes, rolls-back, or reaches the transaction timeout period.

You can roll back the production redeployment process by making the older application version active. This may be necessary if, for example, you determine that there is a problem with the newer version of the application, and you want WebLogic Server to begin moving clients back to the older version. To make the older application version active, redeploy it.

Production Redeployment In Clusters

In a WebLogic Server cluster, each clustered server instance retires its local deployment of the retiring application version when the current workload is completed. This means that an application version may be retired on some clustered server instances before it is retired on other servers in the cluster. Note, however, that in a cluster failover scenario, client failover requests are always handled by the same application version on the secondary server, if the application version is still available. If the same application version is not available on the secondary server, the failover does not succeed.

Requirements and Restrictions for Production Redeployment

In order to use the production redeployment feature, an application must meet certain requirements during the development and deployment phases.

Development Requirements

The production redeployment strategy is supported for:

- Stand-alone Web Application (WAR) modules and Enterprise Applications (EARs) whose clients access the application via a Web application (HTTP).
- Enterprise Applications that are accessed by inbound JMS messages from a global JMS destination, or from inbound JCA requests.
- All types of Web Services, including conversational and reliable Web Services.

Production redeployment is *not* supported for:

- Stand-alone EJB or RAR modules. If you attempt to use production redeployment with such modules, WebLogic Server rejects the redeployment request. To redeploy such modules, remove their version identifiers and explicitly redeploy the modules.
- Applications that use JTS drivers. For more information on JDBC application module limitations, see [JDBC Application Module Limitations](#) in *Configuring and Managing WebLogic JDBC*.

- Applications that obtain JDBC data sources via the DriverManager API; in order to use production redeployment, an application must instead use JNDI to look up data sources.
- Applications that include EJB 1.1 container-managed persistence (CMP) EJBs. To use production redeployment with applications that include CMP EJBs, use EJB 2.x CMP instead of EJB 1.1 CMP.

Production redeployment only supports HTTP clients and RMI clients (see [“Graceful Shut Down of RMI Client Request Processing” on page 8-12](#)). 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.

During development, applications must be designed to meet specific requirements in order to ensure that multiple versions of the application can safely coexist in a WebLogic Server domain during production redeployment. See [Developing Versioned Applications for Production Redeployment](#) in *Developing WebLogic Server Applications* for information about the programming conventions required for applications to use production redeployment.

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*.

WARNING: Because the production redeployment strategy requires an application to observe certain programming conventions, use production redeployment only with applications that are approved by your development and design staff. Using production redeployment with an application that does not follow Oracle’s programming conventions can lead to corruption of global resources or other undesirable application behavior.

Deployment Requirements

A deployed application must specify a version number before you can perform subsequent production redeployment operations on the application. In other words, you cannot deploy a non-versioned application, and later perform production redeployment with a new version of the application.

Restrictions on Production Redeployment Updates

WebLogic Server can host a maximum of two different versions of an application at one time.

When you redeploy a new version of an application, you cannot change:

- An application's deployment targets
- An application's security model
- A Web application's persistent store settings

To change any of the above features, you must first undeploy the active version of the application.

Specifying an Application Version Identifier

WebLogic Server attempts to use the production redeployment strategy when the currently-deployed application and the redeployed application specify different versions.

To assign a version identifier to an application, Oracle recommends that you store a unique version string directly in the `MANIFEST.MF` file of the EAR or WAR being deployed. Your development process should automatically increment the version identifier with each new application release before packaging the application for deployment.

Maintaining version information in this manner ensures that the production redeployment strategy is used with each redeployment of the application in a production domain. See [Developing Versioned Applications for Production Redeployment](#) in *Developing WebLogic Server Applications*.

Assigning a Version Identifier During Deployment and Redeployment

If you are testing the production redeployment feature, or you want to use production redeployment with an application that does not include a version string in the manifest file, specify a unique version string by using the `-appversion` option when deploying or redeploying an application:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -deploy -name myTestDeployment
    -source /myDeployments/myApplication/91Beta
    -targets myCluster -stage -appversion .91Beta
```

The version string specified with `-appversion` is applied only when the deployment source files do not specify a version string in `MANIFEST.MF`. See [Application Version Conventions in *Developing WebLogic Server Applications*](#) for information about valid characters and character formats for application version strings.

WARNING: Do not use `-appversion` to deploy or redeploy in a production environment unless you are certain the application follows Oracle’s programming conventions for production redeployment. See [“Requirements and Restrictions for Production Redeployment” on page 8-6](#). Using production redeployment with an application that does not follow the programming conventions can cause corruption of global resources or other undesirable application behavior.

Displaying Version Information for Deployed Applications

The WebLogic Server Administration Console displays both version string information and state information for all deployed applications and modules. To view this information, select the Deployments node in the left-hand pane of the Administration Console.

You can also display version information for deployed applications from the command line using the `weblogic.Deployer -listapps` command:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -listapps
```

Redeploying a New Version of an Application

To redeploy a new version of an application using the production redeployment strategy:

1. Verify that only one version of the application is currently deployed in the domain. See [“Displaying Version Information for Deployed Applications” on page 8-9](#).
2. Verify that the deployed application and the newer application have different version strings:
 - a. Use the instructions in [“Displaying Version Information for Deployed Applications” on page 8-9](#) to determine the currently-deployed version of the application.
 - b. Examine the version string in the `MANIFEST.MF` file of the new application source you want to deploy:

Redeploying Applications in a Production Environment

```
jar xvf myApp.ear MANIFEST.MF
cat MANIFEST.MF
```

3. Place the new application deployment files in a suitable location for deployment. Oracle recommends that you store each version of an application's deployment files in a unique subdirectory.

For example, if the currently-deployed application's files are stored in:

```
/myDeployments/myApplication/91Beta
```

You would store the updated application files in a new subdirectory, such as:

```
/myDeployments/myApplication/1.0GA
```

WARNING: For `nostage` or `external_stage` mode deployments, do not overwrite or delete the deployment files for the older version of the application. The original deployment files are required if you later choose to roll back the retirement process and revert to the original application version.

4. Redeploy the new application version and specify the updated deployment files. If the updated deployment files contain a unique version identifier in the manifest file, use a command similar to:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -redeploy -name myTestDeployment
    -source /myDeployments/myApplication/1.0GA
```

If the new deployment files do not contain a version identifier in the manifest, see [“Assigning a Version Identifier During Deployment and Redeployment” on page 8-8](#).

By default WebLogic Server makes the newly-redeployed version of the application active for processing new client requests. Existing clients continue to use the older application until their work is complete and the older application can be safely undeployed.

If you want to specify a fixed time period after which the older version of the application is retired (regardless of whether clients finish their work), use the `-retiretimeout` option with the `-redeploy` command:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -redeploy -name myTestDeployment
    -source /myDeployments/myApplication/1.0GA
    -retiretimeout 300
```

`-retiretimeout` specifies the number of seconds after which the older version of the application is retired. You can also retire the older application immediately by using the `-undeploy` command and specifying the older application version, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -undeploy -name myTestDeployment
  -appversion .91Beta
```

5. Verify that both the old and new versions of the application are deployed, and that the correct version of the application is active. See [“Displaying Version Information for Deployed Applications”](#) on page 8-9.

Undeploying a Retiring Application

If WebLogic Server has not yet retired an application version, you can immediately undeploy the application version without waiting for retirement to complete. This may be necessary if, for example, an application remains in the retiring state with only one or two long-running client sessions that you do not want to preserve. To force the undeployment of a retiring version of an application, use the **-undeploy** command and specify the application version:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -undeploy -name myTestDeployment
  -appversion .91Beta
```

Notes: You cannot specify the **-graceful** option to the **-undeploy** command when undeploying an application version that is being retired, or waiting for a retirement timeout to occur.

If you do not explicitly specify an application version with the **-appversion** option, WebLogic Server undeploys the active version and all retired versions of the application. If an older version of the application is not yet retired and you run the **-undeploy** command without specifying the **-appversion** option, WebLogic Server logs a warning message in the server log and does not undeploy the unretired version. To later undeploy such versions of the application, you must run the **-undeploy** command again and specify the application version with the **-appversion** option.

Rolling Back the Production Redeployment Process

Reversing the production redeployment process switches the state of the active and retiring applications and redirects new client connection requests accordingly. Reverting the production redeployment process might be necessary if you detect a problem with a newly-deployed version of an application, and you want to stop clients from accessing it.

To roll back the production redeployment process, issue a second **-redeploy** command and specify the deployment source files for the older version, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -redeploy -name myTestDeployment
    -source /myDeployments/myApplication/91Beta
    -retiretimeout 300
```

If the deployment files do not contain a version identifier in the manifest, see [“Assigning a Version Identifier During Deployment and Redeployment”](#) on page 8-8.

Graceful Shut Down of RMI Client Request Processing

You can use the `-rmigraceperiod` attribute to specify a grace period (in seconds) for processing of RMI client requests when retiring (`-retirement`) or gracefully (`-graceful`) shutting down an application. The work manager of a server instance accepts and schedules RMI calls until the grace period expires without a receiving new RMI client request. If a new RMI client request occurs within the grace period, the grace period is reset and RMI client processing continues until:

- Another RMI request resets the grace period or
- The grace period expires without an RMI client request.

If an RMI client tries to access a retired application or a shutdown application, the client receives `NoSuchObjectException` when trying to invoke the object. If a new version of the object is available, an application needs to catch the `NoSuchObjectException` exception and look up the new version of the object using the JNDI environment property `weblogic.jndi.WLContext.RELAX_VERSION_LOOKUP` in *Developing Applications with WebLogic Server* to return bindings from the active version of the application.

For example, the following command places the application in Administration mode after allowing any pending HTTP sessions or in-process work to complete:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -stop -adminmode -graceful
    -rmigraceperiod 30
```

If graceful retirement is initiated at T_0 seconds, and a RMI request (`msg1`) arrives at T_1 seconds where $T_1 < (T_0 + 30)$, the application waits for $T_1 + 30$ seconds for additional RMI client requests.

- If `msg2` arrives at T_2 seconds, where $T_2 < (T_1 + 30)$, `msg2` is scheduled and the grace period is reset to $T_2 + 30$ seconds.

- If there are no additional requests or a request (msg2) arrives at T2 seconds, where $T2 > (T1 + 30)$, the work manager stops scheduling RMI requests.


Distributing a New Version of a Production Application

When you distribute a new version of an application, WebLogic Server prepares the new application version for deployment. You can then deploy the application in Administration mode, which makes it available only via a configured Administration channel. External clients cannot access an application that has been distributed and deployed in Administration mode.

You can use the `-adminmode` option to start the application in administration mode. For information, see [“Starting a Distributed Application in Administration Mode” on page 6-22](#).

The older version of the application remains active to process both new and existing client requests. WebLogic Server does not automatically retire the older version of the application when you distribute and deploy a newer version in Administration mode.

Figure 8-2 Distributing a New Version of an Application



After you complete testing of the new application via an Administration channel, you either undeploy the new application version or start it. Starting the application causes WebLogic Server to route new client connections to the updated application, and begin retiring the older application version.

Steps for Distributing a New Version of an Application

The basic steps for distributing a new version of an application in Administration mode are nearly the same as those documented in [“Redeploying a New Version of an Application” on page 8-9](#). You simply use the `weblogic.Deployer -distribute` command, rather than the `-redeploy` command, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -distribute -name myTestDeployment
    -source /myDeployments/myApplication/1.0GA
    -appversion 1.0GA
```

Once the application is distributed, start it in Administration mode:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -start -adminmode -name myTestDeployment
    -source /myDeployments/myApplication/1.0GA
    -appversion 1.0GA
```

Starting the application in Administration mode makes it available only via a configured Administration channel. See [Configuring Network Resources](#) in *Designing and Configuring WebLogic Server Environments*.

Optionally, specify the retirement policy or timeout period for distributed applications.

Making an Application Available to Clients

After performing final testing using the configured Administration channel, you can open the distributed version of the application that is running in Administration mode to new client connections by using the `weblogic.Deployer -start` command without the `-adminmode` option:

1. Use the Administration Console to view the version and state information of both application versions:
 - a. Verify that both versions of the application are still deployed.

- b. Note the version identifier of the application version that is running in Administration mode.
2. Use the `-appversion` option to `weblogic.Deployer` to start the application that was distributed and deployed in Administration mode:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -start -name myTestDeployment
    -appversion .91Beta
```

By default WebLogic Server routes new client requests to the application version that was previously distributed and running in Administration mode. Existing clients continue using the older application until their work is complete and the older application can be safely undeployed. If you want to specify a fixed time period after which the older version of the application is retired (regardless of whether clients finish their work), use the `-retiretimeout` option:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -start -name myTestDeployment
    -appversion .91Beta -retiretimeout 300
```

`-retiretimeout` specifies the number of seconds after which the older version of the application is retired.

3. Use the Administration Console to verify that the previously-distributed application is now active, and that the former application version is now retiring. See [“Displaying Version Information for Deployed Applications” on page 8-9](#).

Best Practices for Using Production Redeployment

When using production redeployment, keep in mind the following best practices:

- Never specify a version string for a production application unless you are certain the application follows Oracle’s programming conventions for production redeployment. See [Developing Versioned Applications for Production Redeployment in *Developing WebLogic Server Applications*](#).
- It is easiest to use production redeployment when applications are deployed in stage mode. With stage mode, WebLogic Server automatically creates a separate directory for each different version of the application that is deployed. These directories are stored in the configured staging directory for the server (by default, the `server_name/stage` subdirectory of the domain directory) and are removed when an the associated application version is undeployed.

- If you deploy in `nostage` mode, store each new version of an application in a dedicated subdirectory. This ensures that you do not overwrite older versions of your deployment files, and allows you to revert to an earlier application version if you detect problems after an update.
- If you deploy in `external_stage` mode, you must store the deployment files for each application version in the correct version subdirectory of each target server's staging directory. For example, the versioned application files used in the previous sample commands would need to be copied into the subdirectories `/myTestDeployment/.91Beta` and `/myTestDeployment/1.0GA` in each server's staging directory before deployment and redeployment.

Using In-Place Redeployment for Applications and Stand-alone Modules

In-place redeployment immediately undeploys and replaces a running application's deployment files with updated deployment files. When used to redeploy entire applications or J2EE modules, in-place redeployment makes the application or module unavailable during the deployment process, and can cause existing clients to lose in-process work. This disruption of client services occurs because application class files and libraries are immediately undeployed from their class loaders and then replaced with updated versions.

Because in-place redeployment of applications and modules adversely affects clients of the application, it should not be used with production applications unless:

- No clients are currently using the application, or
- It is acceptable to lose client access to the application and in-process work during redeployment.

WARNING: WebLogic Server always performs in-place redeployment for applications that do not include a version identifier.

WebLogic Server also uses in-place redeployment for many partial redeployment operations (redployment commands that affect only a portion of an application). In some cases, using partial redeployment is acceptable with production applications, because the redeployed files do not adversely affect active client connections. [Table 8-2](#) describes each type of partial deployment and its effect on deployed applications.

Table 8-2 Partial Redeployment Behavior

Scope of Partial Redeployment	Redeployment Behavior	Recommended Usage
Graphics files, static HTML files, JSPs	Source files are immediately replaced on-disk and served on the next client request.	Safe for production applications.
J2EE Modules in an Enterprise Application	All files are immediately replaced. Java class files and libraries are unloaded from class loaders and replaced with updated files.	Use only during scheduled application downtime, or when it is not critical to preserve client connections and in-process work.
Deployment plan with dynamic property changes (such as tuning parameters)	The application is updated in-place. If the application is versioned, the plan version is <i>not</i> incremented.	Safe for all production environments.
Deployment plan with non-dynamic property (resource binding) changes	<p>If the application is versioned, is compatible with production redeployment, and is redeployed, WebLogic Server increments the version identifier and uses the production redeployment strategy to update the application.</p> <p>If the application cannot use production redeployment, you must redeploy the entire application.</p>	<p>Safe for versioned applications that are compatible with production redeployment. See “Using Production Redeployment to Update Applications” on page 8-5.</p> <p>If the application cannot use production redeployment, update the deployment plan only during scheduled application downtime or when it is not critical to preserve client connections and in-process work.</p> <p>You must redeploy (instead of update) applications with deployment plans that contain changes to non-dynamic properties. Attempts to update applications with such plans will fail.</p>

Redeploying Applications and Modules In-Place

To redeploy an entire application or stand-alone module using the in-place redeployment strategy:

1. If you want to preserve client connections to the application, first take the application offline and verify that no clients are accessing the application.

The exact method for taking an application offline depends on the architecture of your WebLogic Server domain. In most cases, a redundant server or cluster is created to host a separate copy of the application, and load balancing hardware or software manages access to both servers or clusters. To take the application offline, the load balancing policies are changed to roll all client connections from one set of servers or clusters to the redundant set.

2. Place the new application deployment files in a suitable location for deployment. Oracle recommends that you store each version of an application's deployment files unique subdirectories.

For example, if the currently deployed application's files (e.g., `SimpleEAR.ear`) are stored in:

```
/myDeployments/myApplication/91Beta/SimpleEAR.ear
```

You would store the updated application files in a new directory, such as:

```
/myDeployments/myApplication/1.0GA/SimpleEAR.ear
```

3. Redeploy the application and specify the updated deployment source files. To redeploy the application on all configured target servers, specify only the deployment name, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic  
-password weblogic -redeploy -source  
/myDeployments/myApplication/1.0GA/SimpleEAR.ear -name myApp
```

If an application was previously deployed to multiple, non-clustered server instances, you can specify a target list to redeploy the application on only a subset of the target servers, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic  
-password weblogic -redeploy -source  
/myDeployments/myApplication/1.0GA/SimpleEAR.ear -name myApp  
-targets myserver1,myserver2
```

Note: For applications deployed to a cluster, redeployment always occurs on all target server instances in the cluster. If the application was previously deployed to all servers in the cluster, you cannot subsequently redeploy the application on a subset of servers in the cluster.

4. If you took the server or cluster hosting the application offline, bring the host servers back online after the redeployment completes.
5. If necessary, restore the load balancing policies of your load balancing hardware or software to migrate clients from the temporary servers back to the online production servers.

Best Practices for Redeploying Applications and Modules In-Place

When using in-place redeployment to redeploy entire applications or stand-alone modules, keep in mind the following:

- Redeploying entire, staged applications may have performance implications due to increased network traffic when deployment files are copied to the Managed Servers. If you have very large applications, consider using the `external_stage` mode and copying deployment files manually before you redeploy the application. See [“Using External_stage Mode Deployment” on page 6-19](#).
- Applications deployed to a WebLogic Server cluster must always be redeployed on all members of the cluster. You cannot redeploy a clustered application to a subset of the cluster.
- To successfully redeploy an Enterprise Application, all of the application’s modules must redeploy successfully.
- By default, WebLogic Server destroys current user sessions when you redeploy a Web application. If you want to preserve Web application user sessions during redeployment, set `save-sessions-enabled` to “true” in the `container-descriptor` stanza of the `weblogic.xml` deployment descriptor file. Note, however, that the application still remains unavailable while in-place redeployment takes place.

Using Partial Redeployment for J2EE Module Updates

The `weblogic.Deployer` utility uses a different command form if you want to redeploy individual modules of a deployed Enterprise Application. To redeploy a subset of the modules of an Enterprise Application, specify `modulename@servername` in the target server list to identify the modules you want to redeploy. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
  -password weblogic -redeploy -name myApp
  -targets mymodule1@myserver1,mymodule2@myserver2
```

Note: The use of `-redeploy module-uri` is deprecated. Instead, use `production redeployment` or `redeploy` the module using the `-targets module@target` syntax.

If the application was previously deployed to a cluster, you must redeploy the module to the entire cluster, rather than a subset of servers. If you specify a subset of servers in the cluster, `weblogic.Deployer` responds with the error:

```
An attempt to add server target target_name to module module_name has been rejected . This is because its parent cluster, cluster_name, is also targeted by the module.
```

Restrictions for Updating J2EE Modules in an EAR

The following restrictions apply to using partial redeployment for modules in an Enterprise Application:

- If redeploying a single J2EE module in an Enterprise Application would affect other J2EE modules loaded in the same class loader, `weblogic.Deployer` requires that you explicitly redeploy all of the affected modules. If you attempt to use partial redeployment with only a subset of the affected J2EE modules, `weblogic.Deployer` displays the error:

```
Exception:weblogic.management.ApplicationException: [J2EE:160076] You must include all of [module_list] in your files list to modify [module]
```

- Remember that if you change an application's deployment descriptor files, the container redeloys the entire application even if you attempt a partial redeployment.
- JAR files in `WEB-INF/lib` cannot be redeployed independently of the rest of the Web application. The container automatically redeloys the entire application, but maintains the state, when redeploying JAR files in `WEB-INF/lib`.

Best Practices for Updating J2EE Modules in an EAR

Keep in mind these best practices when using partial redeployment for Enterprise Application modules:

- When you use partial redeployment to redeploy a J2EE module in an Enterprise Application, all classes loaded in the classloader for the updated module are reloaded. You can define custom class loading hierarchies in the WebLogic Server deployment descriptor to minimize the impact of partial redeployment to other modules in the application. See [WebLogic Server Application Classloading](#) in *Developing WebLogic Server Applications* for more information on class loading behavior.

- Classes in the `WEB-INF/classes` directory can be redeployed independently of the rest of the Web application. You can also deploy only the updated classes (rather than the entire `WEB-INF/classes` directory) by setting the Reload Period for the Web application. (See [weblogic.xml Schema](#) in *Developing Web Applications, Servlets, and JSPs for WebLogic Server* for more information.)
- By default, WebLogic Server destroys current user sessions when you redeploy a Web application module. If you want to preserve Web application user sessions during redeployment, set `save-sessions-enabled` to “true” in the `container-descriptor` stanza of the `weblogic.xml` deployment descriptor file. Note, however, that the application still remains unavailable while in-place redeployment takes place.

Updating Static Files in a Deployed Application

In a production environment, you may occasionally need to refresh the static content of a Web application module—HTML files, image files, JSPs, and so forth—without redeploying the entire application. If you deployed a Web application or an Enterprise Application as an exploded archive directory, you can use the `weblogic.Deployer` utility to update one or more changed static files in-place.

To redeploy a single file associated within a deployment unit, specify the file name at the end of the redeploy command. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name myApp -redeploy myApp/copyright.html
```

Always specify the pathname to updated files relative to the root directory of the exploded archive directory. In the above example, the Web application is deployed as part of an Enterprise Application, so the module directory is specified (`myApp/copyright.html`).

If the Web application module had been deployed as a stand-alone module, rather than as part of an Enterprise Application, the file would have been specified alone (`copyright.html`).

You can also redeploy an entire directory of files by specifying a directory name instead of a single file. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name myApp -redeploy myApp/myjsp
```

In the above example, all files and subdirectories located in the `myjsp` subdirectory of the Enterprise Application are redeployed in-place.

Updating the Deployment Configuration for an Application

After you have deployed an application or stand-alone module, you can change the WebLogic Server deployment configuration using either the Administration Console or `weblogic.Deployer`.

The Administration Console enables you to interactively modify individual deployment configuration properties, while `weblogic.Deployer` only allows you to specify an updated deployment plan file to use for reconfiguring the application.

Modifying a Configuration Using the Administration Console

The Administration Console enables you to reconfigure all deployment configuration properties for an application, including properties that were not included in the application's deployment plan. If an application was deployed with a deployment plan, the Console displays any deployment plan configuration properties in the plan in the Deployment Plan tab for the application.

The remaining configuration tabs for an application enable you to change other WebLogic Server configuration properties. The exact properties that are available for configuration depend on the type of application or J2EE module that is deployed. These tabs are available regardless of whether or not the application was deployed with an deployment plan.

Note that certain configuration changes are safe to apply to running production applications, while other changes require you to shut down and restart the application. See [“Understanding Redeployment Behavior for Deployment Configuration Changes”](#) on page 8-23.

How Configuration Changes Are Stored

When you use the Administration Console to make changes to properties that were defined in a deployment plan, the Console generates a new deployment plan that containing variable definitions for the new properties you modified, as well as any existing variables defined in the plan. You can select where to save the new deployment plan.

Updating an Application to Use a Different Deployment Plan

The `weblogic.Deployer` utility enables you to update an application's deployment configuration by providing a new deployment plan to use with the application.

Note: The updated deployment plan must be valid for the application's current target servers, or the configuration update will fail.

To reconfigure an application with a different (and valid) deployment plan, use the **-update** command and specify the new deployment plan file, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -update -name myTestDeployment
    -plan /myDeployments/myNewPlan.xml
```

Note: Certain configuration changes are safe to apply to running production applications, while other changes require you to shut down and restart the application. See [“Understanding Redeployment Behavior for Deployment Configuration Changes” on page 8-23](#).

Understanding Redeployment Behavior for Deployment Configuration Changes

Note: The redeployment behavior in this section applies both to configuration changes made using the Administration Console and via the `weblogic.Deployer -update` command.

When you change the deployment configuration for a deployed application, WebLogic Server applies the changes using a redeployment operation. The type of redeployment strategy used depends on the nature of configuration changes applied. If you modified the value of a resource binding property, the configuration update uses either the production redeployment strategy (if the application supports production redeployment), or the in-place redeployment strategy for the entire application. Performing in-place redeployment for an entire application or module is not recommended for production environments, because the application is first undeployed, causing connected clients to lose in-progress work. See [“Overview of Redeployment Strategies” on page 8-1](#) for more information.

Redeploying Applications in a Production Environment

Managing Deployed Applications

The following sections describe how to perform common maintenance tasks on applications and modules that are currently deployed to a WebLogic Server domain:

- [“Taking a Production Application Offline”](#) on page 9-1
- [“Undeploying Shared Libraries and Packages”](#) on page 9-3
- [“Adding a New Module to a Deployed Enterprise Application”](#) on page 9-3
- [“Changing the Order of Deployment at Server Startup”](#) on page 9-4
- [“Changing the Target List for an Existing Deployment”](#) on page 9-6
- [“Removing Files from a Web Application Deployment”](#) on page 9-6
- [“Managing Long-Running Deployment Tasks”](#) on page 9-7
- [“On-demand Deployment of Internal Applications”](#) on page 9-8

Taking a Production Application Offline

WebLogic Server provides two different ways to take an application offline for testing or maintenance purposes:

- [Stopping an Application to Restrict Client Access](#)—Makes an application unavailable for processing client requests, but does not remove the deployment from the WebLogic Server domain. Stopping an application places the deployment in Administration mode, which allows you to perform internal testing using a configured Administration channel.

- **Undeploying an Application or Module**—Makes an application unavailable for processing client requests and removes WebLogic Server-generated deployment files from the domain.

Stopping an Application to Restrict Client Access

As described in “[Distributing Applications to a Production Environment](#)” on page 6-21, distributing an application and starting it in Administration mode, restricts access to an application to a configured Administration channel. You can also stop a running application to client requests and place it in Administration mode. In a production environment, you may want to stop an application to confirm a reported problem, or to isolate the application from external client processing in order to perform scheduled maintenance.

Use the **-stop** command to place a running application into Administration mode, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -stop -adminmode
```

By default, WebLogic Server immediately stops the application, without regard to pending HTTP sessions or in-process work. If you want to wait for pending HTTP sessions to complete their work before stopping the application to client requests and placing it in Administration mode, add the **-graceful** option:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -stop -adminmode -graceful
```

Note: If you do not explicitly specify an application version with the **-appversion** option, the **-stop** command will only stop the active version of the application. If there are other versions of the application that you also want to stop (or that you want to stop instead of the active version), you must specify them with the **-appversion** option.

To restart an application that was previously stopped, making it available to external clients, use the **-start** command and specify the deployment name. You do not need to redeploy a stopped application to make it generally available:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -start
```

Undeploying an Application or Module

After you deploy a new application or stand-alone module to servers in a domain, the deployment name remains associated with the deployment files you selected. Even after you stop the deployment on all servers, the files remain available for redeployment using either the Administration Console or `weblogic.Deployer` utility.

If you want to remove a deployment name and its associated deployment files from the domain, you must explicitly undeploy the application or stand-alone module. To undeploy a deployment unit from the domain using `weblogic.Deployer`, specify the `-undeploy` command, as in:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -undeploy
```

Note: Using the `-undeploy` command without the `-targets` and `-submoduletargets` flags completely removes the application or stand-alone module from all WebLogic Server instances and untargets all JMS sub-module resources.

By default, WebLogic Server immediately stops and undeploys the application, interrupting current clients and in-progress work. For a production application, you may want to undeploy “gracefully,” allowing current HTTP clients to complete their work before undeploying:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -undeploy -graceful
```

Undeploying a deployment unit does not remove the original source files used for deployment. It only removes the deployment’s configuration from the domain, as well as any deployment files that WebLogic Server created during deployment (for example, files copied with stage deployment mode and files uploaded to the Administration Server).

If you need to redeploy a deployment unit after undeploying it, you must again identify the deployment files, targets, staging mode, and deployment name using the instructions in [“Deploying Applications and Modules with weblogic.deployer” on page 6-1](#).

Undeploying Shared Libraries and Packages

A shared J2EE libraries or optional package cannot be undeployed until all applications that reference the library or package are first undeployed. If no applications reference an application or package, use the instructions in [“Undeploying an Application or Module” on page 9-2](#) to undeploy it.

Adding a New Module to a Deployed Enterprise Application

WebLogic Server’s module-level targeting support enables you to add and deploy a new Enterprise Application module without having to redeploy other modules that are already deployed. To deploy a new module in an EAR, you simply use module-level targeting syntax described in [“Module-Targeting Syntax” on page 6-10](#).

For example, if you were to add a module, `newmodule.war`, to a deployed Enterprise Application named `myapp.ear` (update `application.xml` file as necessary), you could then deploy `newmodule.war` using the `weblogic.Deployer` command:

```
java weblogic.Deployer -username myname -password mypassword
    -name myapp.ear -deploy -targets newmodule.war@myserver
    -source /myapp/myapp.ear
```

This command deploys the new module without redeploying the other modules in the application. Note that you must specify the correct file extension (`.war` in the above example) for archived modules in an EAR file.

Changing the Order of Deployment at Server Startup

By default, WebLogic Server deploys applications and resources in the following order:

1. JDBC system modules
2. JMS system modules
3. J2EE Libraries and optional packages
4. Applications and stand-alone modules
5. Startup classes

Note: WebLogic Server security services are always initialized before server resources, applications, and startup classes are deployed. For this reason, you cannot configure custom security providers using startup classes, nor can custom security provider implementations rely on deployed server resources such as JDBC.

Changing the Deployment Order for Applications and Stand-alone Modules

You can change the deployment order for a deployed application or *stand-alone* module by setting the `AppDeploymentMBean DeploymentOrder` attribute in the Administration Console (or programmatically using the `AppDeploymentMBean`). The `DeploymentOrder` attribute controls the load order of deployments relative to one another—modules with lower `DeploymentOrder` values deploy before those with higher values. By default, each deployment unit is configured with a `Deployment Order` value of 100. Deployments with the same `Deployment Order` value are deployed in alphabetical order using the deployment name. In all

cases, applications and stand-alone modules are deployed after the WebLogic Server instance has initialized dependent subsystems.

Note: You cannot change the load order of applications and stand-alone modules using the `weblogic.Deployer` utility.

To view or change the deployment order of deployments using the Administration Console, follow the steps in [Change the server deployment order](#) in *Administration Console Online Help*.

Changing the Deployment Order for Modules in an Enterprise Application

The modules contained in an Enterprise Application are deployed in the order in which they are declared in the `application.xml` deployment descriptor. See [Enterprise Application Deployment Descriptor Elements](#) in *Developing WebLogic Server Applications*.

Ordering Startup Class Execution and Deployment

By default WebLogic Server startup classes are run after the server initializes JMS and JDBC services, and after applications and stand-alone modules have been deployed.

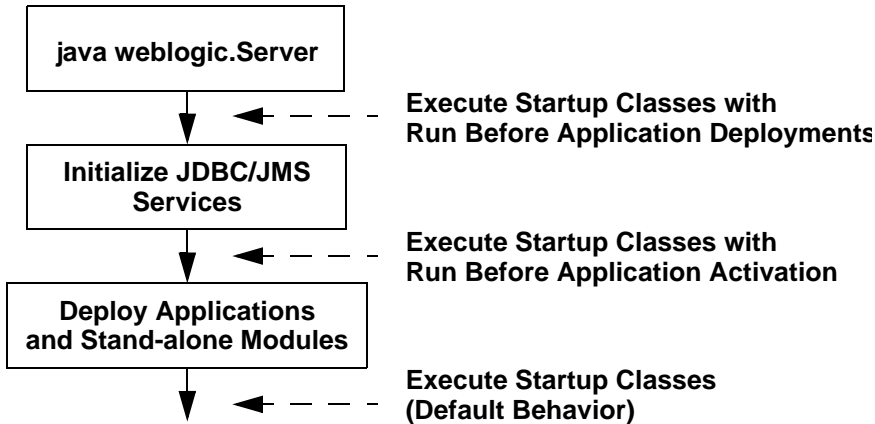
If you want to perform startup tasks after JMS and JDBC services are available, but before applications and modules have been activated, you can select the Run Before Application Deployments option in the Administration Console (or set the `StartupClassMBean`'s `LoadBeforeAppActivation` attribute to “true”).

If you want to perform startup tasks before JMS and JDBC services are available, you can select the Run Before Application Activations option in the Administration Console (or set the `StartupClassMBean`'s `LoadBeforeAppDeployments` attribute to “true”).

To select Run Before Applications or Run Before Application Activations in the Administration Console, see [Configure startup classes](#) in *Administration Console Online Help*.

The following figure summarizes the time at which WebLogic Server executes startup classes.

Figure 9-1 Startup Class Execution



See the full [Javadocs](#) for `StartupClassMBean` for more information.

Changing the Target List for an Existing Deployment

After you deploy an application or stand-alone module in a WebLogic Server domain, you can change the target server list to add new WebLogic Server instances or to remove existing server instances. If you remove a target server, only the target list itself is updated—the deployment unit remains deployed to the removed server until you explicitly undeploy it. Similarly, if you add a new target server, you must explicitly deploy the deployment unit on the new server before it is active on that server.

To add a new server to the target list using `weblogic.Deployer`, simply specify the new list of target servers with the `-deploy` command. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mydeploymentname -deploy
    -targets server1, newserver
```

Removing Files from a Web Application Deployment

If you deploy a Web application using an exploded archive directory, you can update static contents of the Web application either by refreshing the files (see [“Updating Static Files in a](#)

Deployed Application” on page 8-21), or by deleting files from the deployment. To delete files, you must use the `weblogic.Deployer` utility with the `delete_files` option. For example:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mywebapp -delete_files mywebapp/copyright.html
```

Always specify the pathname to updated files starting with the top level of the exploded archive directory. In the above example, the Web application resides in an exploded archive directory named `mywebapp`.

Note: Because the `-delete_files` option deletes all specified files or, if you specify a directory but do not specify files within the directory, all files in the specified directory, Oracle recommends that you use caution when using the `delete_files` option and that you do not use the `delete_files` option in production environments.

Managing Long-Running Deployment Tasks

The WebLogic Server deployment system automatically assigns a unique ID to deployment tasks so that you can track and manage their progress. `weblogic.Deployer` enables you to assign your own task identification numbers for use with deployment commands, as well as monitor and cancel long-running deployment tasks.

For example, the following command assigns a task ID of `redeployPatch2` to a new deployment operation:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -name mymodule -targets myserver -id redeployPatch2
    -nowait -deploy c:\localfiles\myapp.ear
```

You can later check the status of the task:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -id redeployPatch2 -list
```

You can check the status of all running tasks:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -listtask
```

If a task takes too long to complete you can cancel it:

```
java weblogic.Deployer -adminurl http://localhost:7001 -user weblogic
    -password weblogic -id redeployPatch2 -cancel
```

On-demand Deployment of Internal Applications

There are many internal applications that are deployed during startup. These internal applications consume memory and require CPU time during deployment. This contributes to the WLS startup time and base memory footprint. Since many of these internal applications are not needed by every user, WLS has been modified to wait and deploy these applications on the first access (on-demand) instead of always deploying them during server startup. This reduces startup time and memory footprint.

Internal Application Types

There are two different types of internal applications. The first type displays a user interface, and includes the Administration Console, UDDI explorer, and WLS test client. The second type does not display a user interface, and includes UDDI and internal servlets for deployment and management file distribution.

For applications with a user interface, WLS displays a status page indicating that on-demand deployment is in progress. This page refreshes every two seconds. When the internal application completes deployment, you are redirected to the internal application. This status page is displayed on the first access of each application. Subsequent invocations do not deploy the application and go directly to the user interface for the internal application.

Configuring On-Demand Deployment

For a development domain, the default is for WLS to deploy internal applications on demand. For a production-mode domain, the default is for WLS to deploy internal applications as part of server startup. You can control the default behavior by configuring the `InternalAppsDeployOnDemandEnabled` attribute in the Domain MBean. You can change the configuration setting using the Administration Console or by using the WebLogic Scripting Tool (WLST).

Configuring On-Demand Deployment Using the Administration Console

To change the `InternalAppsDeployOnDemandEnabled` attribute using the Administration Console, perform the following steps:

1. Start an edit session with the **Lock and Edit** button.
2. Select the domain to bring up the Configuration -> General tab.

3. Change the setting of the **Enable on-demand deployment of internal applications** check box.
4. Click **Save** and then click **Activate Changes** to activate the changes that will take effect at the next restart of WLS.

Configuring On-Demand Deployment Using WLST

To change the `InternalAppsDeployOnDemandEnabled` attribute using the WLST, perform the following steps:

```
connect()  
edit()  
startEdit()  
cmo.setInternalAppsDeployOnDemandEnabled(false)  
activate()
```

Managing Deployed Applications

weblogic.Deployer Command-Line Reference

`weblogic.Deployer` is a Java-based deployment tool that provides administrators and developers command-line based deployment operations.

Note: See the [WLST Command and Variable Reference](#) for information about performing deployment operations using the WebLogic Scripting Tool (WLST).

The following sections describe the `weblogic.Deployer` utility:

- [“Required Environment for weblogic.Deployer”](#) on page A-1
- [“Syntax for Invoking weblogic.Deployer”](#) on page A-2
 - [“SSL Arguments”](#) on page A-2
 - [“Connection Arguments”](#) on page A-4
 - [“User Credentials Arguments”](#) on page A-4
 - [“Common Arguments”](#) on page A-5
- [“Commands and Options”](#) on page A-7
- [“Example config.xml File and Corresponding weblogic.Deployer Command”](#) on page A-29

Required Environment for weblogic.Deployer

To set up your environment to use the `weblogic.Deployer` utility:

1. Install and configure the WebLogic Server software, as described in the [WebLogic Server Installation Guide](#).

2. Add the WebLogic Server classes to the CLASSPATH environment variable, and ensure that the correct JDK binaries are available in your PATH. You can use the `setWLSEnv.sh` or `setWLSEnv.cmd` script, located in the `server/bin` subdirectory of the WebLogic Server installation directory, to set the environment.
3. If you are connecting to an Administration Server through a configured Administration channel, you must also configure SSL on the machine on which you run `weblogic.Deployer`. See [Configuring iSSL in *Managing WebLogic Security*](#) for instructions about configuring SSL.

Syntax for Invoking weblogic.Deployer

```
java [ SSL Arguments ] weblogic.Deployer [ Connection Arguments ]  
    [ User Credentials Arguments ] COMMAND-NAME command-options  
    [ Common Arguments ]
```

Command names and options are not case-sensitive. See [“Commands and Options” on page A-7](#) for detailed syntax and examples of using `weblogic.Deployer` commands.

SSL Arguments

```
java [ -Dweblogic.security.TrustKeyStore=DemoTrust ]  
    [ -Dweblogic.security.JavaStandardTrustKeystorePassPhrase=password ]  
    [ -Dweblogic.security.CustomTrustKeystoreFileName=filename  
    -Dweblogic.security.TrustKeystoreType=CustomTrust  
    [ -Dweblogic.security.CustomTrustKeystorePassPhrase=password ]  
    ]  
    [ -Dweblogic.security.SSL.hostnameVerifier=classname ]  
    [ -Dweblogic.security.SSL.ignoreHostnameVerification=true ]  
weblogic.Deployer  
    [ User Credentials Arguments ]  
    COMMAND-NAME command-arguments
```

If you have enabled the domain-wide administration port, or if you want to secure your administrative request by using some other listen port that is secured by SSL, you must include

SSL arguments when you invoke `weblogic.Deployer`. [Table A-1](#) describes all SSL arguments for the `weblogic.Deployer` utility.

Table A-1 SSL Arguments

Argument	Definition
<code>-Dweblogic.security.TrustKeyStore=DemoTrust</code>	<p>Causes <code>weblogic.Deployer</code> to trust the CA certificates in the demonstration trust keystore (<code>WL_HOME\server\lib\DemoTrust.jks</code>).</p> <p>This argument is required if the server instance to which you want to connect is using the demonstration identity and certificates.</p> <p>By default, <code>weblogic.Deployer</code> trusts only the CA certificates in the Java Standard Trust keystore (<code>SDK_HOME\jre\lib\security\cacerts</code>).</p>
<code>-Dweblogic.security.JavaStandardTrustKeystorePassPhrase=password</code>	<p>Password that was used to secure the Java Standard Trust keystore.</p> <p>If the Java Standard Trust keystore is protected by a password, and if you want to trust its CA certificates, you must use this argument.</p> <p>By default, the Java Standard Trust keystore is not protected by a password.</p>
<code>-Dweblogic.security.CustomTrustKeyStoreFileName=filename</code> <code>-Dweblogic.security.TrustKeystoreType=CustomTrust</code>	<p>Causes <code>weblogic.Deployer</code> to trust the CA certificates in a custom keystore that is located at <i>filename</i>. You must use both arguments to trust custom keystores. The <i>filename</i> must match exactly the <code>ServerMBean.CustomTrustKeyStoreFileName</code> value persisted in <code>config.xml</code>; if the value specified in the <code>CustomTrustKeyStoreFileName</code> attribute is a relative pathname, you must also specify the same relative pathname in this argument.</p>
<code>-Dweblogic.security.CustomTrustKeystorePassPhrase=password</code>	<p>Password that was used to secure the custom keystore.</p> <p>You must use this argument only if the custom keystore is protected by a password.</p>
<code>-Dweblogic.security.SSL.hostnameVerifier=classname</code>	<p>Name of a custom Host Name Verifier class. The class must implement the <code>weblogic.security.SSL.HostnameVerifier</code> interface.</p>
<code>-Dweblogic.security.SSL.ignoreHostnameVerification=true</code>	<p>Disables host name verification.</p>

Connection Arguments

```
java [SSL Arguments] weblogic.Deployer
    [-adminurl protocol://listen_address:port_number]
    [User Credentials Arguments] COMMAND-NAME command-options [Common
Arguments]
```

Most `weblogic.Deployer` commands require you to specify the `-adminurl` arguments described in [Table A-2](#) to connect to an Administration Server instance.

Table A-2 Connection Arguments

Argument	Definition
<code>-adminurl</code> <code>[protocol://]Admin-Server-listen-address:listen-port</code>	<p>Listen address and listen port of the Administration Server.</p> <p>To use a port that is not secured by SSL, the format is <code>-adminurl [protocol]Admin-Server-listen-address:port</code> where <code>t3</code>, <code>http</code>, <code>iiop</code>, and <code>iiops</code> are valid protocols.</p> <p>In order to use an <code>adminurl</code> with the HTTP protocol, you must enable the HTTP tunneling option in the Administration Console. For more information, see Setting Up WebLogic Server for HTTP Tunneling in <i>Designing and Configuring WebLogic Server Environments</i>. For instructions on enabling HTTP tunneling in the Administration Console, see Configure HTTP protocol in <i>Administration Console Online Help</i>.</p> <p>To use a port that is secured by SSL, the format is <code>-adminurl secure-protocol://Admin-Server-listen-address:port</code> where <code>t3s</code> and <code>https</code> are valid secure protocols.</p> <p>To connect to the Administration Server via a configured Administration channel, you must specify a valid administration port number: <code>-adminurl secure-protocol://Admin-Server-listen-address:domain-wide-admin-port</code></p> <p>There is no default value for this argument.</p>

User Credentials Arguments

```
java [ SSL Arguments ] weblogic.Deployer [Connection Arguments]
    [ { -username username [-password password] } |
    [ -userconfigfile config-file [-userkeyfile admin-key] ] ]
    COMMAND-NAME command-options [Common Arguments]
```


Most `weblogic.Deployer` commands require you to provide the user credentials of a WebLogic Server administrator.

Table A-3 User Credentials Arguments

Argument	Definition
<code>-username <i>username</i></code>	Administrator username. If you supply the <code>-username</code> option but you do not supply a corresponding <code>-password</code> option, <code>weblogic.Deployer</code> prompts you for the password.
<code>-password <i>password</i></code>	<p>Password of the Administrator user.</p> <p>To avoid having the plain text password appear in scripts or in process utilities such as <code>ps</code>, first store the username and encrypted password in a configuration file using the <code>STOREUSERCONFIG</code> command with <code>weblogic.Admin</code>. Omit both the <code>-username</code> and <code>-password</code> options to <code>weblogic.Deployer</code> to use the values stored in the default configuration file. Before specifying the <code>-password <i>password</i></code> attribute, you must first generate the file using the WebLogic Scripting Tool (WLST) <code>storeUserConfig</code> command as described in the WLST Command and Variable Reference.</p> <p>If you want to use a specific configuration file and key file, rather than the default files, use the <code>-userconfigfile</code> and <code>-userkeyfile</code> options to <code>weblogic.Deployer</code>.</p>
<code>-userconfigfile <i>config-file</i></code>	Location of a user configuration file to use for the administrative username and password. Use this option, instead of the <code>-user</code> and <code>-password</code> options, in automated scripts or in situations where you do not want to have the password shown on-screen or in process-level utilities such as <code>ps</code> . Before specifying the <code>-userconfigfile</code> attribute, you must first generate the file using the WebLogic Scripting Tool (WLST) <code>storeUserConfig</code> command as described in the WLST Command and Variable Reference .
<code>-userkeyfile <i>admin-key</i></code>	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 <code>-userconfigfile</code> option). Before specifying the <code>-userkeyfile</code> attribute, you must first generate the file using the WebLogic Scripting Tool (WLST) <code>storeUserConfig</code> command as described in the WLST Command and Variable Reference .

Common Arguments

```
java [ SSL Arguments ] weblogic.Deployer [ Connection Arguments ]
    [ { -username username [-password password] } ] |
```

```
[ -userconfigfile config-file [-userkeyfile admin-key] ] ]  
COMMAND-NAME command-options [Common Arguments]
```

The common arguments described in [Table A-4](#) can be used with any commands described in “[Commands and Options](#)” on page A-7.

Table A-4 Common Arguments for weblogic.Deployer

Argument Name	Description
-advanced	Prints full command-line help text for all <code>weblogic.Deployer</code> actions and options.
-debug	Display debug messages in the standard output.
-examples	Display example command lines for common tasks.
-help	Prints command-line help text for the most commonly used <code>weblogic.Deployer</code> actions and options.
-noexit	By default <code>weblogic.Deployer</code> calls <code>System.exit(1)</code> if an exception is raised during command processing. The exit value displayed indicates the number of failures that occurred during the deployment operation. The <code>-noexit</code> option overrides this behavior for batch processing.
-noversion	Indicates that <code>weblogic.Deployer</code> 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 <code>external_stage</code> staging mode. If you use this option, you cannot use versioned applications.
-nowait	<code>weblogic.Deployer</code> prints the task ID and exits without waiting for the action to complete. This option initiates multiple tasks and then monitor them later with the <code>-list</code> action.
-output <raw formatted>	(Deprecated.) Either <code>raw</code> or <code>formatted</code> to control the appearance of <code>weblogic.Deployer</code> output messages. Both output types contain the same information, but <code>raw</code> output does not contain embedded tabs. By default, <code>weblogic.Deployer</code> displays <code>raw</code> output.

Table A-4 Common Arguments for `weblogic.Deployer`

Argument Name	Description
<code>-purgetasks</code>	Indicates that <code>weblogic.Deployer</code> should flush out deployment tasks that are retired.
<code>-remote</code>	Indicates that <code>weblogic.Deployer</code> is not running on the same machine as the Administration Server, and that source paths specified in the command are valid for the Administration Server machine itself. If you do not use the <code>-remote</code> option, <code>weblogic.Deployer</code> assumes that all source paths are valid paths on the local machine.
<code>-timeout <i>seconds</i></code>	Maximum time, in seconds, to wait for the deployment task to complete. After the time expires, <code>weblogic.Deployer</code> prints out the current status of the deployment and exits.
<code>-verbose</code>	Additional progress messages, including details about the prepare and activate phases of the deployment.
<code>-version</code>	Prints version information for <code>weblogic.Deployer</code> .

Commands and Options

The following sections describe the `weblogic.Deployer` commands and command options used to perform deployment tasks with WebLogic Server:

- [Cancel](#)
- [Deploy](#)
- [Distribute](#)
- [Listapps](#)
- [List, Listtask](#)
- [Redeploy](#)
- [Start](#)
- [Stop](#)
- [Undeploy](#)

- [Update](#)

Note: `weblogic.Deployer` commands are displayed in **bold** type to distinguish them from command options.

Cancel

Attempt to cancel a running deployment task.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -cancel task_id
    [Common Arguments]
```

Command or Option	Definition
<i>task_id</i>	Identifier of the deployment task to cancel. Specify the identifier by using the <code>id</code> option with the <code>deploy</code> , <code>distribute</code> , <code>update</code> , <code>undeploy</code> , <code>redeploy</code> , <code>stop</code> , and <code>start</code> commands.

Examples

The following command starts a deployment operation and specifies the task identifier, `myDeployment`:

```
java weblogic.Deployer -adminurl http://localhost:7001
    -username weblogic -password weblogic
    -deploy ./myapp.ear -id myDeployment
```

If the deployment task has not yet completed, the following command attempts to cancel the deployment operation:

```
java weblogic.Deployer -adminurl http://localhost:7001
    -username weblogic -password weblogic
    -cancel -id myDeployment
```

Deploy

Deploys or redeploys an application or module.

Note: The `-ACTIVATE` command, an alias for `deploy`, is deprecated.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -deploy [[-name] deployment_name] [-source] file
    [-plan file] [-targets target_list] [-submoduletargets target_list]
    [-upload]
    [-stage | -nostage | -external_stage]
    [-retiretimeout seconds]
    [-library [-libspectver version] [-libimplver version]]
    [-altappdd file] [-altwlsappdd file]
    [-securityModel] [-enableSecurityValidation]
    [-id task_id]
    [Common Arguments]
```

Command or Option	Definition
<code>-name deployment_name</code>	<p>Deployment name to assign to a newly-deployed application or stand-alone module.</p> <p>Both the <code>-name</code> option and <code>deployment_name</code> argument are optional, as described in the Syntax. If a deployment name is not explicitly identified with the <code>deploy</code> command, the name is derived from the specified deployment file or directory:</p> <ul style="list-style-type: none"> For an archive file, the default deployment name is the full name of the archive file with the file extension. For example, when deploying <code>myear.ear</code>, the default deployment name is <code>myear.ear</code>. For an exploded archive directory, the default deployment name is the name of the top-level directory. If you specify an application installation root directory, the default deployment name is derived from the archive filename or exploded archive directory name in the <code>/app</code> subdirectory.
<code>-source file</code>	<p>Archive file or exploded archive directory to deploy. You can omit the <code>-source</code> option and supply only the file or directory to deploy.</p>

Command or Option	Definition (Continued)
<code>-plan file</code>	Deployment plan to use when deploying the application or module. By default, <code>weblogic.Deployer</code> does not use an available deployment plan, even if you are deploying from an application root directory that contains a plan.
<code>-targets target_list</code>	<p>Targets on which to distribute and deploy the application or module. The <code>target_list</code> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<code><module1>@<server1></code>). This enables you to deploy different modules of an Enterprise Application to different servers or clusters.</p> <p>If you do not specify a target list with the deploy command, the target defaults to:</p> <ul style="list-style-type: none"> • the Administration Server instance for new deployments. • the application's current targets for deployed applications.
<code>-submoduletargets target_list</code>	JMS Server targets for resources defined within a JMS application module. See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i> .
<code>-upload</code>	Transfers the specified deployment files, including deployment plans and alternate deployment descriptors, to the Administration Server. Use this option when you are on a remote machine and you cannot copy the deployment files to the Administration Server by other means. The application files are uploaded to the WebLogic Server Administration Server's upload directory prior to distribution and deployment.

Command or Option	Definition (Continued)
-stage -nostage -external_stage	<p>Staging mode to use when deploying or distributing an application:</p> <ul style="list-style-type: none"> • -stage—Copies deployment files to target servers' staging directories. <code>stage</code> is the default mode used when deploying or distributing to Managed Server targets. • -nostage—Does not copy the deployment files to target servers, but leaves them in a fixed location, specified by the <code>-source</code> option. Target servers access the same, copy of the deployment files. <code>nostage</code> is the default used when deploying or distributing to the Administration Server (for example, in a single-server domain). • -external_stage—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 manually copy the files or use a third-party tool or script. <p>See “Controlling Deployment File Copying with Staging Modes” on page 6-14.</p>
-retiretimeout <i>seconds</i>	<p>Number of seconds before WebLogic Server retires the currently-running version of this application or module. See “Redeploying a New Version of an Application” on page 8-9.</p>
-library	<p>The deployment as a shared J2EE library or optional package. You must include the <code>-library</code> option when deploying or distributing any J2EE library or optional package. See “Deploying Shared Java EE Libraries and Dependent Applications” on page 6-22.</p>
-libspeclver <i>version</i>	<p>Specification version of a J2EE library or optional package. This option can be used only if the library or package does not include a specification version in its manifest file. <code>-libversion</code> can be used only in combination with <code>-library</code>. See “Registering Libraries with WebLogic Server” on page 6-24.</p>
-libimplver <i>version</i>	<p>Implementation version of a J2EE library or optional package. This option can be used only if the library or package does not include an implementation version in its manifest file. <code>-libimplversion</code> can be used only in combination with <code>-library</code>. See “Registering Libraries with WebLogic Server” on page 6-24.</p>

Command or Option	Definition (Continued)
<code>-usenonexclusivelock</code>	Deployment operation uses an existing lock, already acquired by the same user, on the domain. This option is helpful in environments where multiple deployment tools are used simultaneously and one of the tools has already acquired a lock on the domain configuration.
<code>-altappdd <i>file</i></code>	(Deprecated.) Name of an alternate J2EE deployment descriptor (<code>application.xml</code>) to use for deployment.
<code>-altwlsappdd <i>file</i></code>	(Deprecated.) Name of an alternate WebLogic Server deployment descriptor (<code>weblogic-application.xml</code>) to use for deployment.
<code>-securityModel</code> [<code>DDOnly</code> <code>CustomRoles</code> <code>CustomRolesAndPolicy</code> <code>Advanced</code>]	Security model to use for this deployment.
<code>-enableSecurityValidation</code>	Enable validation of security data.
<code>-id <i>task_id</i></code>	Task identifier of a running deployment task. You can specify an identifier with the distribute , deploy , redploy , start , or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.

Examples

See the following sections for examples of using the **-deploy** command:

- [“Deploying to a Single-Server Domain” on page 6-4](#)
- [“Deploying an Application with a Deployment Plan” on page 6-4](#)
- [“Uploading Deployment Files from a Remote Client” on page 6-3](#)
- [“Deploying to One or More Targets” on page 6-7](#)
- [“Deploying to a Cluster Target” on page 6-8](#)
- [“Using Module-Level Targeting for Deploying an Enterprise Application” on page 6-9](#)
- [“Targeting Application-Scoped JMS, JDBC, and WLDF Modules” on page 6-11](#)
- [“Using Sub-Module Targeting with JMS Application Modules” on page 6-12](#)

- “Using Nostage Mode Deployment” on page 6-17
- “Using Stage Mode Deployment” on page 6-18
- “Using External_stage Mode Deployment” on page 6-19
- “Distributing Applications to a Production Environment” on page 6-21
- “Registering Libraries with WebLogic Server” on page 6-24
- “Deploying Applications That Reference Libraries” on page 6-25

Distribute

Prepares deployment files for deployment by copying deployment files to target servers and validating them.

A distributed application can be started quickly with the [Start](#) command. You can start the application in Administration mode, or make it available to Administration and client requests. While in Administration mode, the application can be accessed only by internal clients through a configured Administration port. External clients cannot access the application.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -distribute [[-name] deployment_name] [-source] file
    [-plan file] [-targets target_list] [-submoduletargets target_list]
    [-upload]
    [-stage | -nostage | -external_stage]
    [-library [-libspectver version] [-libimplver version]]
    [-altappdd file] [-altwlsappdd file]
    [-securityModel] [-enableSecurityValidation]
    [-id task_id]
    [Common Arguments]
```

Command or Option	Definition
<code>-name <i>deployment_name</i></code>	<p>Deployment name to assign to the distributed application or module.</p> <p>Both the <code>-name</code> option and <code>deployment_name</code> argument are optional, as described in the Syntax. If a deployment name is not explicitly identified, a name is derived from the specified deployment file or directory:</p> <ul style="list-style-type: none"> • For an archive file, the default deployment name is the name of the archive file without the file extension (<code>myear</code> for the file <code>myear.ear</code>). • For an exploded archive directory, the default deployment name is the name of the top-level directory. • If you specify an application installation root directory, the default deployment name is derived from the archive filename or exploded archive directory name in the <code>/app</code> subdirectory.
<code>-source <i>file</i></code>	Archive file or exploded archive directory to distribute. You can omit the <code>-source</code> option and supply only the file or directory.
<code>-plan <i>file</i></code>	Deployment plan to distribute with the application or module, used to configure the application.
<code>-targets <i>target_list</i></code>	<p>Targets on which to distribute the application or module.</p> <p>The <code>target_list</code> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<code><module1>@<server1></code>). This enables you to distribute different modules of an Enterprise Application to different servers or clusters.</p> <p>If you do not specify a target list with the deploy command, the target defaults to:</p> <ul style="list-style-type: none"> • The Administration Server instance for new deployments. • The application's current targets for deployed applications.
<code>-submoduletargets <i>target_list</i></code>	JMS Server targets for resources defined within a JMS application module. See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i> .

Command or Option	Definition (Continued)
<code>-upload</code>	Transfers the specified deployment files, including any specified deployment plans, to the Administration Server before distribution. Use this option when you are on a remote machine and you cannot copy the deployment files to the Administration Server by other means. The application files are uploaded to the WebLogic Server Administration Server's upload directory prior to distribution.
<code>-stage -nostage -external_stage</code>	<p>Staging mode to use when deploying or distributing an application:</p> <ul style="list-style-type: none"> • <code>-stage</code>—Copies deployment files to target servers' staging directories. <code>stage</code> is the default mode used when deploying or distributing to Managed Server targets. • <code>-nostage</code>—Does not copy the deployment files to target servers, but leaves them in a fixed location, specified by the <code>-source</code> option. Target servers access the same copy of the deployment files. <code>nostage</code> is the default used when deploying or distributing to the Administration Server (for example, in a single-server domain). • <code>-external_stage</code>—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 manually copy the files or use a third-party tool or script. <p>See “Controlling Deployment File Copying with Staging Modes” on page 6-14.</p>
<code>-library</code>	Identifies the deployment as a shared J2EE library or optional package. You must include the <code>-library</code> option when deploying or distributing any J2EE library or optional package. See “Deploying Shared Java EE Libraries and Dependent Applications” on page 6-22.
<code>-libspever <i>version</i></code>	Specification version of a J2EE library or optional package. This option can be used only if the library or package does not include a specification version in its manifest file. <code>-libversion</code> can be used only in combination with <code>-library</code> . See “Registering Libraries with WebLogic Server” on page 6-24.
<code>-libimplver <i>version</i></code>	Implementation version of a J2EE library or optional package. This option can be used only if the library or package does not include an implementation version in its manifest file. <code>-libimplversion</code> can be used only in combination with <code>-library</code> . See “Registering Libraries with WebLogic Server” on page 6-24.

Command or Option	Definition (Continued)
<code>-altappdd file</code>	(Deprecated.) Name of an alternate J2EE deployment descriptor (<code>application.xml</code>) to use for deployment or distribution.
<code>-altwlsappdd file</code>	(Deprecated.) Name of an alternate WebLogic Server deployment descriptor (<code>weblogic-application.xml</code>) to use for deployment or distribution.
<code>-securityModel</code> [<code>DDOnly</code> <code>CustomRoles</code> <code>CustomRolesAndPolicy</code> <code>Advanced</code>]	Security model to be used for this application.
<code>-enableSecurityValidation</code>	Enable validation of security data.
<code>-id task_id</code>	Task identifier of a running deployment task. You can specify an identifier with the distribute , deploy , redploy , start , or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.

Examples

The **distribute** command operates similar to **deploy**, but WebLogic Server does not start the application or module on target servers. See the examples links for [“Deploy” on page A-9](#) command for more information.

Listapps

Lists the deployment names for applications and stand-alone modules deployed, distributed, or installed to the domain.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -listapps
    [Common Arguments]
```

Command or Option	Definition
<code>-listapps</code>	Lists the deployment names for applications and stand-alone modules deployed, distributed, or installed to the domain

Examples

See [“Displaying Version Information for Deployed Applications”](#) on page 8-9.

List, Listtask

Displays the status of deployment tasks currently running in the domain.

Syntax

```
java [SSL Arguments] weblogic.Deployer Connection Arguments
      [User Credentials Arguments] <-list | -listtask> [task_id]
      [Common Arguments]
```

Command or Option	Definition
<code>task_id</code>	Identifier of a deployment task to display.

Examples

See [“Managing Long-Running Deployment Tasks”](#) on page 9-7.

Redeploy

Redeploys a running application or part of a running application.

Syntax

```
java [SSL Arguments] weblogic.Deployer
      Connection Arguments [User Credentials Arguments]
      -redeploy [[-name] deployment_name] {-source file | filelist}
      [-plan file] [-targets target_list] [-submoduletargets target_list]
      [-upload]
```

```
[-delete_files]
[-retiretimeout seconds] [-id task_id]
[-rmiGracePeriod seconds]
[Common Arguments]
```

Command or Option	Definition
<code>-name <i>deployment_name</i></code>	Deployment name of a deployed application or module. The <code>-name</code> option can be omitted, in which case the name is taken from the <code>-source <i>file</i></code> argument.
<code>-source <i>file</i></code>	<p>Archive file or exploded archive directory to distribute, deploy, or redeploy.</p> <p>When used with the <code>redeploy</code> command, the <code>-source</code> option specifies the location of new deployment files to redeploy, for example, when updating an application to a new version.</p> <p>To specify multiple files for a partial redeployment, omit the <code>-source</code> option and supply only a <i>filelist</i>.</p> <p>Note: To redeploy an entire J2EE module within an Enterprise Application, use the module-targeting syntax, <code>-targets <i>module@target</i></code>, described in “Using Partial Redeployment for J2EE Module Updates” on page 8-19.</p>
<i>filelist</i>	<p>One or more files to redeploy. If the <i>filelist</i> specifies multiple files, the redeployment is treated as a partial redeployment of the specified files.</p> <p>Note: Use a <i>filelist</i> specification only for redeploying static files <i>within</i> a J2EE module. To redeploy an entire J2EE module within an Enterprise Application, use the module-targeting syntax, <code>-targets <i>module@target</i></code>, described in “Using Partial Redeployment for J2EE Module Updates” on page 8-19.</p> <p>The use of <code>-redeploy <i>module-uri</i></code> is deprecated. Instead, use production redeployment or redeploy the module using the <code>-targets <i>module@target</i></code> syntax.</p>

Command or Option	Definition (Continued)
<code>-plan file</code>	<p>Deployment plan to use when distributing, deploying, or redeploying.</p> <p>When redeploying an application, the <code>-plan</code> option allows you to specify an updated configuration to use during the redeployment. If the revised deployment plan contains changes to resource bindings, WebLogic Server attempts to redeploy a new version of the application alongside an older version. See “Updating the Deployment Configuration for an Application” on page 8-22.</p>
<code>-targets target_list</code>	<p>Targets on which to distribute, deploy, or redeploy the application or module.</p> <p>The <code>target_list</code> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<code><module1>@<server1></code>). This enables you to redeploy different modules of an Enterprise Application to different servers or clusters.</p> <p>If you do not specify a target list with the deploy command, the target defaults to:</p> <ul style="list-style-type: none"> • The Administration Server instance for new deployments. • The application’s current targets for deployed applications. <p>If you do not specify a target list with the redeploy command, the application is redeployed on all of its current target servers.</p>
<code>-submoduletargets target_list</code>	<p>JMS Server targets for resources defined within a JMS application module. See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i>.</p>
<code>-upload</code>	<p>Transfers the specified deployment files, including deployment plans and alternate deployment descriptors, to the Administration Server. Use this option when you are on a remote machine and you cannot copy the deployment files to the Administration Server by other means. The application files are uploaded to the WebLogic Server Administration Server’s upload directory prior to distribution and deployment.</p>

Command or Option	Definition (Continued)
<code>-delete_files</code>	<p>Removes static files from a server's staging directory. <code>delete_files</code> is valid only for unarchived deployments, and only for applications deployed using <code>-stage</code> mode. You must specify target servers when using this option, as shown in the following example:</p> <pre>java weblogic.Deployer -adminurl http://myserver:7001 -username weblogic -password weblogic -name myapp -targets myapp@myserver -redploy -delete_files myapp/tempindex.html</pre> <p><code>delete_files</code> only removes files that WebLogic Server copied to the staging area during deployment. If you use the <code>delete_files</code> option with an application that was deployed using either <code>-nostage</code> or <code>-external_stage</code> mode, the command does not delete the files.</p> <p><code>delete_files</code> can only be used in combination with the redploy command.</p> <p>Note: Because the <code>-delete_files</code> option deletes all specified files or, if you specify a directory but do not specify files within the directory, all files in the specified directory, Oracle recommends that you use caution when using the <code>delete_files</code> option and that you do not use the <code>delete_files</code> option in production environments.</p>
<code>-retiretimeout <i>seconds</i></code>	<p>Number of seconds before WebLogic Server retires the currently-running version of this application or module. See “Redeploying a New Version of an Application” on page 8-9.</p>
<code>-rmiGracePeriod <i>seconds</i></code>	<p>The amount of time, in seconds, that the work manager accepts and schedules RMI calls until there are no more RMI requests arriving within the RMI grace period during a graceful shutdown or a retirement.</p>
<code>-id <i>task_id</i></code>	<p>Task identifier of a running deployment task. You can specify an identifier with the deploy, redploy, or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.</p>

Examples

See the following sections for examples of using the **redploy** command:

- “Redeploying a New Version of an Application” on page 8-9
- “Rolling Back the Production Redeployment Process” on page 8-11
- “Steps for Distributing a New Version of an Application” on page 8-14
- “Redeploying Applications and Modules In-Place” on page 8-18
- “Using Partial Redeployment for J2EE Module Updates” on page 8-19
- “Updating Static Files in a Deployed Application” on page 8-21

Start

Makes a stopped (inactive) application available to clients on target servers. The **start** command does not redistribute deployment files to target servers. Optionally, with the **adminmode** option, starts the application in Administration mode, which makes it available only via a configured Administration channel. In order to issue a **start** command, the files must already be available through a **deploy** or **distribute** command.

Note: The **activate** command, an alias for **start**, is deprecated.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -start [-adminmode] [-name] deployment_name
    [-appversion version] [-planversion version]
    [-targets target_list] [-submoduletargets target_list]
    [-retiretimeout seconds]
    [-id task_id]
    [Common Arguments]
```

Command or Option	Definition
-adminmode	Start application in Administration mode, not Production mode (which is the default).
-name <i>deployment_name</i>	Deployment name of a deployed application or module. The name option can be omitted, in which case the name is taken directly from the <code>deployment_name</code> . (If the <code>deployment_name</code> specifies a file or directory name, the deployment name is derived from the file specification.)
-appversion <i>version</i>	Version of the application to start.
-planversion <i>version</i>	Version of the deployment plan to use when starting the application.
-targets <i>target_list</i>	<p>Targets on which to DISTRIBUTE, DEPLOY, REDEPLOY, or START the application or module.</p> <p>The <i>target_list</i> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<module1>@<server1>). This enables you to deploy different modules of an Enterprise Application to different servers or clusters.</p> <p>If you do not specify a target list with the deploy command, the target defaults to:</p> <ul style="list-style-type: none"> • The Administration Server instance for new deployments. • The application's current targets for deployed applications. <p>If you do not specify a target list with the redeploy or start commands, the command is performed on all of the application's current targets.</p>
-submoduletargets <i>target_list</i>	JMS Server targets for resources defined within a JMS application module. See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i> .

Command or Option	Definition (Continued)
<code>-retiretimeout <i>seconds</i></code>	Number of seconds before WebLogic Server retires the currently-running version of this application or module. See “Redeploying a New Version of an Application” on page 8-9 .
<code>-id <i>task_id</i></code>	Task identifier of a running deployment task. You can specify an identifier with the distribute , deploy , redeploy , start , or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.

Examples

See the following sections for examples of using the **start** command:

- [“Starting a Distributed Application” on page 6-22](#)
- [“Making an Application Available to Clients” on page 8-14](#)
- [“Stopping an Application to Restrict Client Access” on page 9-2](#)

Stop

Makes an application inactive and unavailable administration and client requests. All of the application’s staged files remain available on target servers for subsequent **start**, **deploy**, **redeploy**, or **undeploy** actions. Optionally, choose to make the application unavailable to client requests by placing it in Administration mode with the **adminmode** option. While in Administration mode, the application be accessed only through a configured Administration channel.

Note: The **deactivate** command, an alias for **stop**, is deprecated.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -stop [-adminmode] [-name] deployment_name
    [-appversion version] [-planversion version]
    [-targets target_list] [-submoduletargets target_list]
    [-ignoresessions] [-graceful] [-rmiGracePeriod seconds]
```

`[-id task_id]`
[\[Common Arguments\]](#)

Argument or Option	Definition
<code>-adminmode</code>	Indicates that a running application should switch to Administration mode and accept only Administration requests via a configured Administration channel. If this option is not specified, the running application is stopped and cannot accept Administration or client requests until it is restarted.
<code>-name <i>deployment_name</i></code>	Deployment name of a deployed application or module. The name option can be omitted, in which case the name is taken directly from the <code>deployment_name</code> . (If the <code>deployment_name</code> specifies a file or directory name, the deployment name is derived from the file specification.)
<code>-appversion <i>version</i></code>	Version identifier of the deployed application.
<code>-planversion <i>version</i></code>	Version identifier of the deployment plan.
<code>-targets <i>target_list</i></code>	<p>Targets on which to distribute, deploy, redploy, start, or stop the application or module.</p> <p>The <i>target_list</i> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<module1>@<server1>). This enables you to deploy different modules of an Enterprise Application to different servers or clusters.</p> <p>If you do not specify a target list with the deploy command, the target defaults to:</p> <ul style="list-style-type: none"> • The Administration Server instance for new deployments. • The application’s current targets for deployed applications. <p>If you do not specify a target list with the redploy, start, or stop commands, the command is performed on all of the application’s current targets.</p>
<code>-submoduletargets <i>target_list</i></code>	JMS Server targets for resources defined within a JMS application module. See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i> .

Argument or Option	Definition (Continued)
<code>-graceful</code>	Stops the application after existing HTTP clients have completed their work. If you do not specify the <code>-graceful</code> option, WebLogic Server immediately stops the application or module. See “Taking a Production Application Offline” on page 9-1.
<code>-rmiGracePeriod <i>seconds</i></code>	The amount of time, in seconds, that the work manager accepts and schedules RMI calls until there are no more RMI requests arriving within the RMI grace period during a <code>graceful</code> shutdown or a <code>retirement</code> .
<code>-ignoresessions</code>	Immediately places the application into Administration mode without waiting for current HTTP sessions to complete.
<code>-id <i>task_id</i></code>	Task identifier of a running deployment task. You can specify an identifier with the distribute , deploy , redploy , start , stop , or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.

Examples

See [“Stopping an Application to Restrict Client Access”](#) on page 9-2.

Undeploy

Stops the deployment unit and removes staged files from target servers.

Note: The **REMOVE** command, an alias for **undeploy**, is deprecated.

WARNING: When you undeploy an application that contains application-scoped resources, the resources are deleted along with the application, which can potentially cause abandoned transactions or lost messages as a result of deleted JMS destinations. For more information, see [Unregister Resource Grace Period](#) in *Programming WebLogic JTA*.

You should only undeploy applications that you are certain you want to completely remove; to temporarily stop client access to applications, use the **stop** command, described in [“Stop”](#) on page A-23, instead.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -undeploy [-name] deployment_name
    [-appversion version] [-planversion version]
    [-targets target_list] [-submoduletargets target_list]
    [-graceful] [-ignoresessions] [-rmiGracePeriod seconds]
    [-id task_id]
    [Common Arguments]
```

Command or Option	Definition
<code>-name <i>deployment_name</i></code>	Deployment name of a deployed application or module. The name option can be omitted, in which case the name is taken directly from the <code>deployment_name</code> . (If the <code>deployment_name</code> specifies a file or directory name, the deployment name is derived from the file specification.)
<code>-appversion <i>version</i></code>	Version identifier of the deployed application.
<code>-planversion <i>version</i></code>	Version identifier of the deployment plan.
<code>-targets <i>target_list</i></code>	<p>Targets from which the application or module are undeployed.</p> <p>Note: Any target not included in the target list is <i>not</i> removed.</p> <p>The <code>target_list</code> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<code><module1>@<server1></code>). This enables you to undeploy different modules of an Enterprise Application from different servers or clusters.</p>
<code>-submoduletargets <i>target_list</i></code>	<p>JMS resources to be undeployed.</p> <p>Note: Any sub-module target not included in the target list is <i>not</i> removed.</p> <p>See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i>.</p>

Command or Option	Definition (Continued)
<code>-graceful</code>	Stops the application after existing HTTP clients have completed their work. If you do not specify the <code>-graceful</code> option, WebLogic Server immediately stops the application or module. See “Taking a Production Application Offline” on page 9-1. The module is undeployed after it is stopped.
<code>-rmiGracePeriod <i>seconds</i></code>	The amount of time, in seconds, that the work manager accepts and schedules RMI calls until there are no more RMI requests arriving within the RMI grace period during a <code>graceful</code> shutdown or a retirement.
<code>-ignoreSessions</code>	Immediately stops and undeploys the application without waiting for current HTTP sessions to complete.
<code>-id <i>task_id</i></code>	Task identifier of a running deployment task. You can specify an identifier with the distributed , deploy , redeploy , start , stop , or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.

Examples

See the following sections for examples of using the **undeploy** command:

- [“Undeploying an Application or Module”](#) on page 9-2
- [“Sub-module Targeting for Stand-alone JMS Modules”](#) on page 6-13

Update

Updates an application’s deployment plan by redistributing the plan files and reconfiguring the application based on the new plan contents.

Note: **update** cannot be used to update an application’s resource bindings. To update the resource bindings for an application, you must use the [Redeploy](#) command.

Syntax

```
java [SSL Arguments] weblogic.Deployer
    Connection Arguments [User Credentials Arguments]
    -update -plan deployment_plan [-name] deployment_name
```

```
[-appversion version] [-planversion version]
[-targets target_list] [-submoduletargets target_list]
[-upload] [-id task_id]
[Common Arguments]
```

Argument or Option	Definition
<code>-plan <i>deployment_plan</i></code>	<p>Deployment plan to use for updating the application's configuration. The specified deployment plan must be valid for the application's target servers. For example, the plan cannot contain null variables for required resources unless those resources were previously defined in the associated descriptors.</p> <p>Update operations update only those descriptors for which there is a changed, not null value in the deployment plan. If a plan that is used by an update operation contains null variables, the current values in the corresponding descriptors are not updated.</p>
<code>-name <i>deployment_name</i></code>	<p>Deployment name of a deployed application or module. The name option can be omitted, in which case the name is taken directly from the <code>deployment_name</code>. (If the <code>deployment_name</code> specifies a file or directory name, the deployment name is derived from the file specification.)</p>
<code>-appversion <i>version</i></code>	Version identifier of the deployed application.
<code>-planversion <i>version</i></code>	Version identifier of the deployment plan.
<code>-targets <i>target_list</i></code>	<p>Targets on which to distribute, deploy, redeploy, undeploy, start, or stop the application or module.</p> <p>The <code>target_list</code> argument is a comma-separated list of the target servers, clusters, or virtual hosts. Each target may be qualified with a J2EE module name (<module1>@<server1>). This enables you to deploy different modules of an Enterprise Application to different servers or clusters.</p> <p>If you do not specify a target list with the deploy command, the target defaults to:</p> <ul style="list-style-type: none"> • The Administration Server instance for new deployments. • The application's current targets for deployed applications. <p>If you do not specify a target list with the redeploy, undeploy, start, or stop commands, the command is performed on all of the application's current targets.</p>

Argument or Option	Definition (Continued)
-submoduletargets <i>target_list</i>	JMS Server targets for resources defined within a JMS application module. See “Using Sub-Module Targeting with JMS Application Modules” on page 6-12 and Using WLST to Manage JMS Servers and JMS System Resources in <i>Configuring and Managing WebLogic JMS</i> .
-upload	Uploads a new deployment plan to the Administration Server before updating the application.
-id <i>task_id</i>	Task identifier of a running deployment task. You can specify an identifier with the distribute , deploy , redeploy , update , start , stop , or undeploy commands, and use it later as an argument to the cancel or list commands. Make sure that the identifier is unique to all other running deployment tasks. The system automatically generates a unique identifier if you do not specify one.

Example

See [“Updating an Application to Use a Different Deployment Plan”](#) on page 8-22.

Example config.xml File and Corresponding weblogic.Deployer Command

This section demonstrates an application’s `config.xml` file and the corresponding `weblogic.Deployer` command to deploy the application.

Assuming:

- `mycluster` is a cluster name
- `D1C2S1` and `D1C2S2` are server names
- `RemoteJMSServer1` and `RemoteJMSServer2` are JMS server names

The application’s `config.xml` file would contain:

```
<AppDeployment
Name="dd-remote-cluster"
SourcePath=". /udd-debug-deployment-on-remote-cluster-jms.xml "
Targets="mycluster">
<SubDeployment Name="RemoteCluster" Targets="mycluster"/>
```

weblogic.Deployer Command-Line Reference

```
<SubDeployment Name="D1C2S2" Targets="D1C2S2" />
<SubDeployment Name="RemoteClusterServers" Targets="D1C2S1,D1C2S2" />
<SubDeployment Name="RemoteClusterJMSServers"
Targets="RemoteJMSServer1,RemoteJMSServer2" />
<SubDeployment Name="RemoteQueue1" Targets="RemoteJMSServer1" />
</AppDeployment>
```

The `weblogic.Deployer -deploy` command to deploy the application would be:

```
java weblogic.Deployer -adminurl t3://MySystem:10000 -username system
-password system -name dd-remote-cluster
-deploy "config\jms\udd-debug-deployment-on-remote-cluster-jms.xml"
-targets mycluster -submoduletargets RemoteCluster@mycluster,
D1C2S2@D1C2S2, RemoteClusterServers@D1C2S1,
RemoteClusterServers@D1C2S2, RemoteClusterJMSServers@RemoteJMSServer1,
RemoteClusterJMSServers@RemoteJMSServer2, RemoteQueue1@RemoteJMSServer1
```

weblogic.PlanGenerator Command Line Reference

`weblogic.PlanGenerator` is a Java-based deployment configuration tool intended for developers who want to export portions of a WebLogic Server deployment configuration into a deployment plan.

Note: See also [WebLogic Scripting Tool](#) for information about performing deployment configuration operations using the WebLogic Scripting Tool (WLST).

The following sections describe how to use the `weblogic.PlanGenerator` utility:

- [“Overview of weblogic.PlanGenerator” on page B-1](#)
- [“Required Environment for weblogic.PlanGenerator” on page B-2](#)
- [“Syntax for Invoking weblogic.PlanGenerator” on page B-2](#)
- [“Common weblogic.PlanGenerator Tasks” on page B-4](#)

Overview of weblogic.PlanGenerator

`weblogic.PlanGenerator` generates WebLogic Server deployment configuration files for an application or stand-alone module. `weblogic.PlanGenerator` provides two primary functions:

- Exporting different categories of WebLogic Server deployment descriptor properties into empty (null) variables in a template deployment plan. You can optionally use an existing deployment plan as input to the configuration export session. Template plans are generally modified further before they can be used. See [“Exporting an Application for Deployment to New Environments” on page 5-1](#).

- Generating a simple initial deployment plan from a J2EE application. See [“Generating a Template Deployment Plan using weblogic.PlanGenerator” on page 5-5](#).

By default, `weblogic.PlanGenerator` writes an application’s deployment plan to a file named `plan.xml` in the application’s root directory. If your application is not in an application root directory, `weblogic.PlanGenerator` writes `plan.xml` to

`<your_dir>/config/deployments/<user>/<application_name>/plan` where:

- `your_dir` is your `TEMP` directory if it is specified in the property `java.io.tmpdir`. If the property `java.io.tmpdir` is not specified, `your_dir` is the WebLogic Server domain directory.
- `user` is your user name.
- `application` is the name of the application.

Required Environment for weblogic.PlanGenerator

To set up your environment to use the `weblogic.PlanGenerator` utility:

1. Install and configure the WebLogic Server software, as described in the [WebLogic Server Installation Guide](#).
2. Add the WebLogic Server classes to the `CLASSPATH` environment variable, and ensure that the correct JDK binaries are available in your `PATH`. You can use the `setWLSEnv.sh` or `setWLSEnv.cmd` script, located in the `server/bin` subdirectory of the WebLogic Server installation directory, to set the environment.

Syntax for Invoking weblogic.PlanGenerator

```
java weblogic.PlanGenerator [Options] [filespec]
```

The `filespec` can be either:

- An absolute or relative path to an archive file
- An absolute or relative path to an exploded archive directory

Oracle recommends:

- That you store your applications in an installation root directory.
- That you specify an application’s installation root directory, with the `-root` option as described in [“Options” on page B-3](#) when you issue a `weblogic.PlanGenerator` command.

In all cases, the application identified with the *filespec* must contain valid J2EE deployment descriptor files.

If you do not specify an application root directory with the `-root` option or a deployment plan path and name with the `-plan` option, by default, `weblogic.PlanGenerator` writes an application's deployment plan to a file named `plan.xml` in the application's root directory. If it cannot locate an application root directory, `weblogic.PlanGenerator` writes `plan.xml` to `<your_dir>/config/deployments/<user>/<application_name>/plan` where:

- *your_dir* is your `TEMP` directory if it is specified in the property `java.io.tmpdir`. If the property `java.io.tmpdir` is not specified, *your_dir* is the WebLogic Server domain directory.
- *user* is your user name.
- *application* is the name of the application.

Options

The following table describes each `weblogic.PlanGenerator` option.

Table B-1 `weblogic.PlanGenerator` Options

Option	Description
<code>-debug</code>	Enables debug mode.
<code>-plan plan_file</code>	Identifies the path and name of the plan file to create for the configuration session.
<code>-useplan plan_file</code>	Existing deployment plan file to initialize from. If you use <code>-root</code> to specify an application root directory, <code>weblogic.PlanGenerator</code> uses the <code>/plan/plan.xml</code> file in the root directory as input, if one is available; otherwise, <code>weblogic.PlanGenerator</code> creates a plan in the root directory. <code>weblogic.PlanGenerator</code> adds additional exported properties to the input plan; any exported properties that were already in the input plan are retained.
<code>-root root_directory</code>	Application root directory on which to perform the plan generation or export.

Table B-1 weblogic.PlanGenerator Options

Option	Description
<p><i>category</i> where valid values for <i>category</i> are:</p> <ul style="list-style-type: none"> • -all • -configurables • -dependencies • -declarations • -dynamics • -none 	<p>Generates null variable definitions in a template deployment plan for different categories of deployment configuration property:</p> <ul style="list-style-type: none"> • all—Creates a plan that exports all editable properties. • configurables—Creates a plan that exports all editable properties except dependencies and declarations. • dependencies—Creates a plan that exports all WebLogic Server descriptor properties that resolve external resource references. This is the default value. • declarations—Creates a plan that exports all properties that declare a resource to other applications and modules. • dynamics—Creates a plan that exports all properties that can be changed on-the-fly without requiring the application to be redeployed. • none—Creates a plan that exports no properties.
<p>-variables [global unique]</p>	<p>Specifies whether variable names created in the deployment plan can be used across all modules of an application, or only within a particular module. For example, a role assignment might be applied to an entire EAR, or to only a single Web application or other module within the EAR. By default, PlanGenerator creates global variables that apply to the entire application.</p>

Common weblogic.PlanGenerator Tasks

The following sections describe common configuration and export tasks, with examples of weblogic.PlanGenerator syntax.

Creating an Initial Deployment Plan in an Application's Root Directory

If you store your application using an installation root directory, and you specify the root directory with the `-root` option, generated deployment plan files are automatically stored in the root directory's `plan` subdirectory.

```
java weblogic.PlanGenerator -root /appRelease/MyApplication
```

In the above example, the `plan.xml` file is automatically stored in `/appRelease/MyApplication/plan`.

Creating a New Deployment Plan Based on an Existing Plan

The following command uses an existing plan as input and generates a new plan in the `/plan` subdirectory of the application root directory:

```
java weblogic.PlanGenerator -useplan /plans/MyApplication_template.xml  
-root /appRelease/MyApplication
```

Controlling the Components Exported to a Deployment Plan

You can use the `-all`, `-configurables`, `-dependencies`, `-declarations`, `-dynamics`, and `-none` options to specify the WebLogic Server deployment descriptor components that are exported to a template deployment plan. The following command exports all configurable properties to null variables in a template deployment plan:

```
java weblogic.PlanGenerator -root /appRelease/MyApplication -all
```

See [“Exporting an Application for Deployment to New Environments”](#) on page 5-1 for more information about exporting a deployment configuration.

weblogic.PlanGenerator Command Line Reference