Oracle® Fusion Middleware

Creating Domains Using the Configuration Wizard

11*g* Release 1 (10.3.1)

E14140-01

May 2009

This document describes how to use the Configuration Wizard to create, update, and extend WebLogic Server domains.

ORACLE®

Oracle Fusion Middleware Creating Domains Using the Configuration Wizard, 11*g* Release 1 (10.3.1)

E14140-01

Primary Author: Kumar Dhanagopal

Contents

5 Customizing JDBC Data Sources and Component Schema

6 Extending WebLogic Domains

List of Figures

List of Tables

Preface

This preface describes the document accessibility features and conventions used in this guide–*Oracle Fusion Middleware Creating Domains Using the Configuration Wizard*.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at `http://www.oracle.com/accessibility/`.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at `http://www.fcc.gov/cgb/consumerfacts/trs.html`, and a list of phone numbers is available at `http://www.fcc.gov/cgb/dro/trsphonebk.html`.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1 Introduction

This guide provides information about configuring WebLogic Server domains by using the Configuration Wizard.

1.1 Introduction to Domains

A domain is the basic administrative unit of WebLogic Server. It consists of one or more WebLogic Server instances, and logically related resources and services that are managed, collectively, as one unit.

Figure 1–1 WebLogic Domain Structure

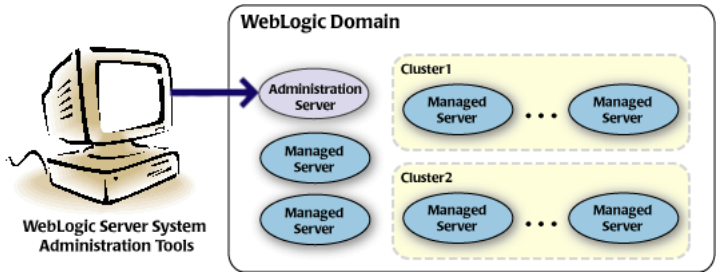

As shown in Figure 1–1, the basic domain infrastructure consists of one administration server and optional managed servers and clusters.

The components of a domain are described in Table 1–1.

Table 1–1 Domain Infrastructure Components

Feature	Description
Administration Server	A domain includes one WebLogic Server instance that is configured as an administration server. All changes to configuration and deployment of applications are done through the administration server. The administration server provides a central point for managing the domain and providing access to the WebLogic Server administration tools. These tools include the following: ▪ WebLogic Server Administration Console: Graphical user interface to the administration server ▪ WebLogic Server Node Manager: A Java program that enables you to start and stop server instances—both administration servers and managed servers—remotely, and to monitor and automatically restart them after an unexpected failure. Note that the node manager is installed on all the machines that host any server instance – both administration server and managed servers.

Table 1–1 (Cont.) Domain Infrastructure Components

Feature	Description
Managed Servers	All other WebLogic Server instances in a domain are called managed servers. Managed servers host application components and resources, which are also deployed and managed as part of the domain. In a domain with only a single WebLogic Server instance, that single server works as both the administration server and the managed server.
Clusters	A domain may also include WebLogic Server clusters, which are groups of managed server instances that work together to provide scalability and high availability for applications. Clusters can improve performance and provide failover when a server instance become unavailable. The servers within a cluster can either run on the same machine or reside in different machines. To the client, a cluster appears as a single WebLogic Server instance.

Note: All managed servers in a domain must run the same version of WebLogic Server. The administration server can run either the same version as the managed servers in the domain, or a later service pack.

In addition to infrastructure components, a domain defines the basic network configuration for the server instances that it contains. Specifically, a domain defines application deployments, supported application services (such as database and messaging services), security options, and physical host machines.

Domain configuration information is stored in the configuration directories under the domain directory.

Common Domain Configurations

You might find it useful to configure multiple domains based on specific criteria, such as system administrator responsibilities, the logical classification of applications, the geographical locations of servers, or size. The following table outlines the most common domain configurations.

Table 1–2 Common Domain Configurations

Configuration	Description
Domain with managed servers	In typical production environments, several managed servers can host applications, and an administration server performs management operations.
Domain with managed servers and clusters	In production environments that require increased performance, throughput, or availability for an application, several managed servers might be grouped in a cluster. In such a case, the domain consists of one or more clusters with the applications they host, additional managed servers (if necessary), and an administration server to perform management operations.
Standalone server domain	In development or test environments, a single application server might be deployed independently without managed servers. In such a case, you can have a domain consisting of a single administration server that also hosts the applications you want to test or develop.

Note: In production environments, it is recommended that you deploy applications only on managed servers; the administration server should be reserved for management tasks.

1.2 Overview of the Configuration Wizard

Before you can develop and run a WebLogic application, you must first create a domain. The Configuration Wizard (illustrated in Figure 1–2), simplifies the process of creating and extending a domain.

To create or extend a domain by using the Configuration Wizard, you simply select the product components to be included in the domain (or choose a template that best meets your requirements), and provide basic configuration information. The Configuration Wizard then creates or extends the domain by using the settings from templates.

For more information about templates, see Section 1.3, "About Domain and Extension Templates."

Figure 1–2 Configuration Wizard

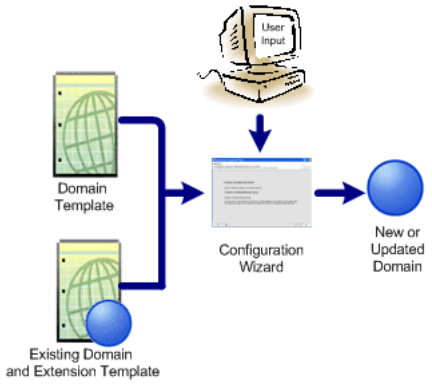

After you create a domain by using the Configuration Wizard, you can start a WebLogic Server instance to run in the domain for developing, testing, and deploying applications.

Note: The Domain Template Builder tool simplifies the process of creating templates by guiding you through the process of creating custom domain and extension templates. You can use these templates for creating and extending domains by using the Configuration Wizard or the WebLogic Scripting Tool (WLST). For information about the Domain Template Builder, see *Oracle WebLogic Server Creating Templates Using the Domain Template Builder*. For information about WLST, see *Oracle WebLogic Scripting Tool*.

Modes of Operation

The Configuration Wizard can be used offline only; that is, when there is no server running. It supports the following modes of operation:

- Graphical mode: An interactive, GUI-based mode
- Console mode: An interactive, text-based mode

> **Note:** For a scripted, silent-mode method, you can use WLST. For more information, see *Oracle WebLogic Scripting Tool*

Output of the Configuration Wizard

A domain created using the Configuration Wizard has the following directories:

- `autodeploy`: This directory provides a location from which you can deploy applications quickly on a development server. When the WebLogic Server instance is running in development mode, it automatically deploys any applications or modules that you place in this directory.
- `bin`: This directory contains scripts to start and stop the administration server, and, optionally, managed servers.
- `config`: This directory contains the following:
  - A domain-specific configuration file, `config.xml`, which specifies the name of the domain and the configuration parameter settings for each server instance, cluster, resource, and service in the domain.
  - Subdirectories that contain the configuration for various system modules: `deployments`, `diagnostics`, `jdbc`, `jms`, `lib`, `nodemanager`, and `security`. These subdirectories contain configuration files that are incorporated, by reference, into the `config.xml` file.

  > **Note:** Depending on your configuration, some subdirectories might not exist.

- `console-ext`—contains console extensions used by the administration server.
- `init-info`: This directory contains files used by the Configuration Wizard to support creation and extension of the domain.
- `lib`: This directory contains the domain library. When the server starts, any `jar` files that you place in this directory are added, dynamically, to the end of the server classpath.
- `security`: This directory contains common security files for all the servers in the domain.
- `servers`: This directory contains a subdirectory for each server in the domain. These server subdirectories, in turn, contain subdirectories that hold directories and files specific to each server in a domain, such as `bin`, `cache`, `data`, `logs`, `security`, and `tmp`.
- `user_staged_config`: If the domain is configured to be user-staged – that is, the administrator is responsible for staging (copying) the configuration information to the managed servers, this directory provides an alternative to the `config` directory.

If the template used to create a domain includes applications, the application files are located, by default, in `user_projects/applications/`*`domain_name`*.

1.3 About Domain and Extension Templates

In the context of the Configuration Wizard, the term **template** refers to a Java Archive (JAR) file, which contains the files and scripts required to create or extend a domain.

The types of templates that can be used by the Configuration Wizard to create or extend domains include:

- **Domain template**: This type of template defines the full set of resources within a domain, including infrastructure components, applications, services, security options, and general environment and operating system options. You can create this type of template from an existing domain by using the Domain Template Builder tool or the `pack` command. Subsequently, you can create a domain based on the template by using the Configuration Wizard.

  The product distribution includes a base WebLogic Server domain template. This template defines the core set of resources within a domain, including an administration server and basic configuration information, infrastructure components, and general environment and operating system options. It does not include sample applications. You can use this template to create a basic WebLogic Server domain, which you can then extend with applications and services, or additional product components.

- **Extension template**: This type of template defines applications and services that can provide additional features, such as Apache Beehive, product sample applications, or JDBC or JMS components. This type of template can be used to update an existing domain.

- **Managed server template**: This type of template defines the subset (within a domain) of resources that are required to create a managed server domain on a remote machine. This type of template can be created by using the `pack` command.

The product installation includes a set of predefined domain and extension templates. This set includes the base WebLogic Server domain template and various extension templates that allow you to add component features and samples to the base domain. For more information about these templates and how they relate to each other, see *Oracle WebLogic Server Domain Template Reference*.

1.4 Creating a Domain

The Configuration Wizard guides you through the process of creating a domain for your target environment by selecting the product components to include in your domain, or by using domain templates. If required, you can also customize the domain to suit your environment by adding and configuring managed servers, clusters, and machine definitions, or customizing predefined JDBC data sources, and JMS file store directories.

You might want to customize your domain in the following circumstances:

- To create a multiserver or clustered domain when using the default settings. All the predefined templates delivered with the product create single-server domains.
- To use a database that is different from the default database in the domain or extension template. In this case, you must customize the JDBC settings to point to the appropriate database.
- To customize the listen port and the SSL port
- To create a test environment by using a domain template that you received, and to modify the domain configuration to work in the test environment based on your requirement

Figure 1–3 summarizes the steps for creating a domain by using the Configuration Wizard. The steps indicated in boxes with dashed lines are relevant only when you

create domains that contain specific products, such as those that use RCU and those that require JDBC data sources.

Figure 1–3 Creating a Domain

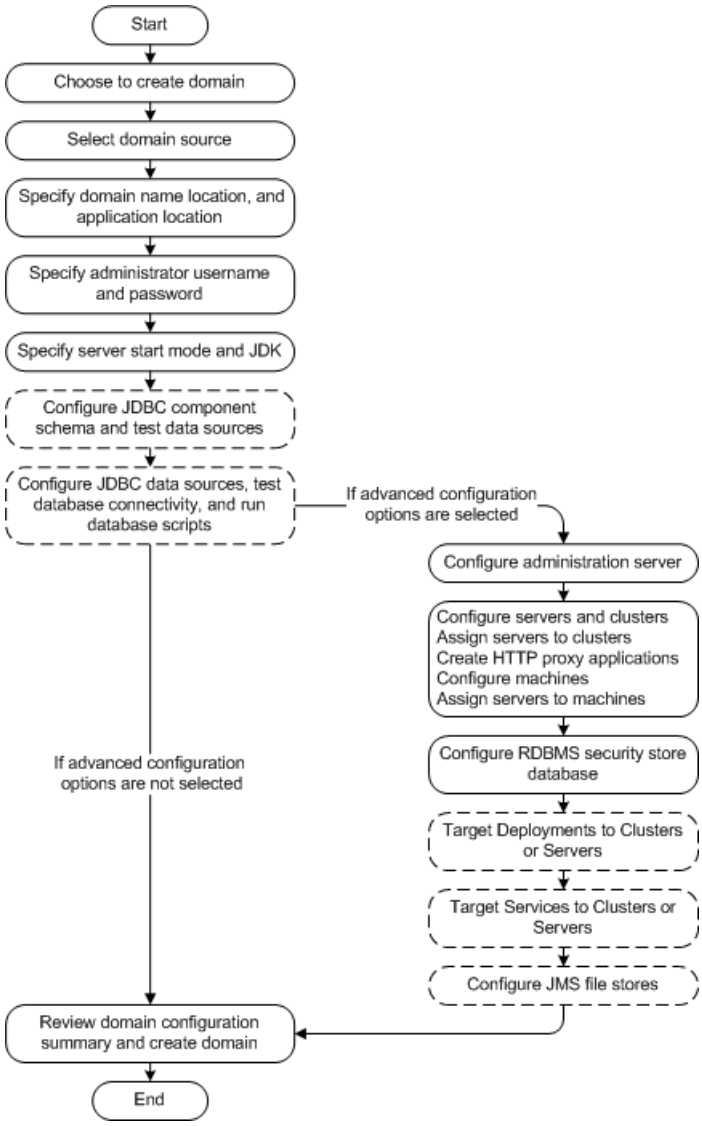

1.5 Extending a Domain

You can extend an existing domain by adding predefined applications and services, or additional product components.

To extend a domain by using the Configuration Wizard, select the domain to extend and then select the additional product component. Alternatively, you can extend an existing domain by specifying an extension template to include additional applications and services. You can also customize the JDBC connections and change the JMS file store. The Configuration Wizard uses your input to update the configuration files, such as `config.xml`, and all other generated components in the domain directory, as required.

Figure 1–4 summarizes the steps for extending a domain by using the Configuration Wizard. The steps indicated in boxes with dashed lines are relevant only when you

extend domains to include specific products, such as those that use RCU and those that require JDBC data sources.

Figure 1–4 Extending a Domain

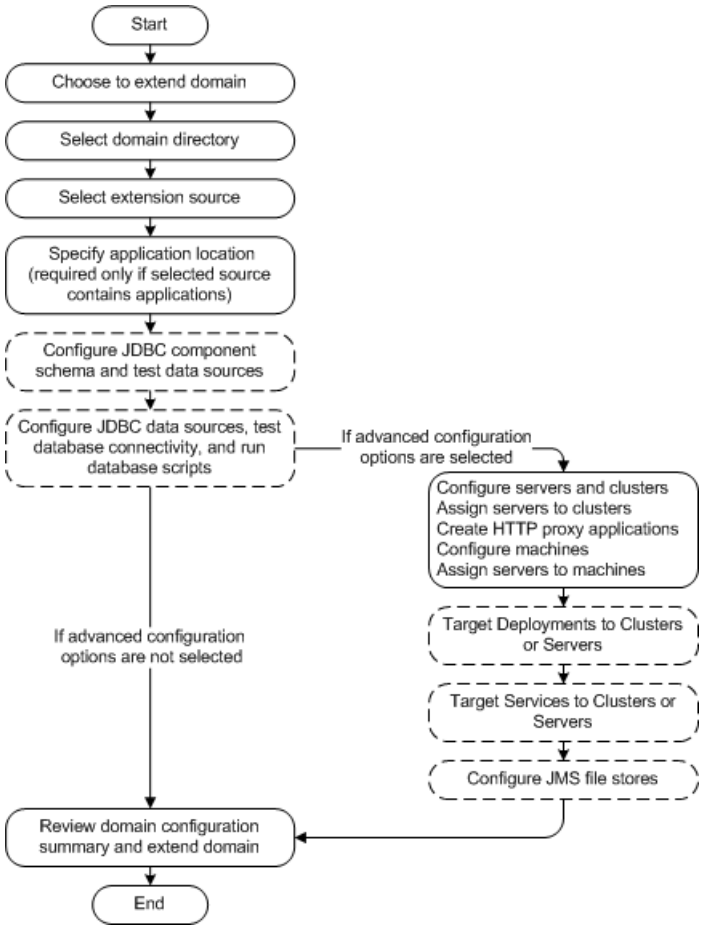

1.6 Additional Tools for Creating, Extending, and Managing Domains

You can create and extend domains by using the Configuration Wizard. In addition, you can use the tools listed in Table 1–3 to create, extend, and manage domains. You can also perform run-time configuration by using the consoles of the product components.

Table 1–3 Additional Tools for Creating, Extending, and Managing Domains

To do this	Use the following tools
Create a domain or extend an existing domain	■ WebLogic Scripting Tool (WLST) WLST is a command-line scripting interface, which you can use to interact with and configure WebLogic Server instances and domains. When WLST is offline, it enables you to create a domain or update an existing domain without connecting to a running WebLogic Server, supporting the same functionality as the Configuration Wizard. For more information, see *WebLogic Scripting Tool*. ■ `unpack` command You can use this command to create a domain from the command line, by using a template that is compatible with your current installation. You cannot use `unpack` to extend an existing domain. For more information, see *Creating Templates and Domains Using the Pack and Unpack Commands*.
Add applications and services, or modify existing settings	■ WebLogic Server administration console For more information, see the *WebLogic Server Administration Console Online Help*. ■ Other system administration tools, such as WLST, weblogic.Admin, JMX, and Ant.
Manage and monitor the health and status of the domain	■ WebLogic diagnostic framework For more information, see *Configuring and Using the WebLogic Diagnostics Framework*. ■ WebLogic Server administration console For more information, see the *WebLogic Server Administration Console Online Help*. ■ WebLogic Server node manager For more information, see *Managing Server Startup and Shutdown*.

2

Starting the Configuration Wizard

This section describes how to start the Configuration Wizard in graphical and console modes.

- Section 2.1, "Starting the Configuration Wizard in Graphical Mode"
- Section 2.2, "Starting the Configuration Wizard in Console Mode"

2.1 Starting the Configuration Wizard in Graphical Mode

The console for the machine on which the product installation resides must support Java-based GUIs. All Windows-based consoles support Java-based GUIs; only a subset of UNIX-based consoles support Java-based GUIs.

> **Note:** If you attempt to start the Configuration Wizard in graphical mode on a system that cannot support graphical display, the Configuration Wizard automatically starts in console mode.

You can start the Configuration Wizard in graphical mode from either the Windows **Start** menu or from the command line.

- To start the Configuration Wizard in graphical mode on a Windows platform, choose **Start** > **Programs** > **Oracle WebLogic** > **WebLogic Server** > **Tools** > **Configuration Wizard**.
- To start the Configuration Wizard in graphical mode from a Windows command prompt or on a UNIX platform:
  1. Log in to the system on which the product is installed.
  2. Open an MS-DOS command prompt window (on Windows) or a command shell (on UNIX).
  3. Go to the `\common\bin` subdirectory of the product installation directory.
  4. Execute the following command:

     On Windows: `config.cmd`

     On UNIX: `sh config.sh`

Note: When you run the `config.cmd` or `config.sh` command, the following error message might be displayed:

```
*sys-package-mgr*: can't create package cache dir
```

The error message indicates that the default cache directory is not valid.

You can change the cache directory by including the `-Dpython.cachedir=<valid_directory>` option in the command line.

The **Welcome** screen is displayed.

2.2 Starting the Configuration Wizard in Console Mode

When run in console mode, the Configuration Wizard is executed in a text-based environment.

To start the Configuration Wizard in console mode:

1. Log in to the system on which the product installation resides.
2. Open an MS-DOS command prompt window (on Windows) or a command shell (on UNIX).
3. Go to the `\common\bin` subdirectory of the product installation directory.
4. Execute the following command:
   - Windows: `config.cmd -mode=console`
   - UNIX: `sh config.sh -mode=console`

   The command and arguments must be entered in lowercase.

   Note: When you run the `config.cmd` or `config.sh` command, the following error message might be displayed:

   ```
   *sys-package-mgr*: can't create package cache dir
   ```

   The error message indicates that the default cache directory is not valid.

   You can change the cache directory by including the `-Dpython.cachedir=<valid_directory>` option in the command line.

The **Welcome** text is displayed.

To proceed, respond to the prompts by entering the number associated with your choice, pressing the **Enter** key. The arrow (`->`) adjacent to a choice indicates the current selection.

- To accept the current selection, type `next` (or `n`) and press the **Enter** key.
- To close the Configuration Wizard, enter `exit` (or `x`) and press the **Enter** key.
- To review or change earlier selections, enter `previous` (or `p`) and press the **Enter** key.

3

Creating a WebLogic Domain

This chapter describes the procedure to create a WebLogic domain by using the Configuration Wizard in graphical mode.

1. Start the Configuration Wizard as described in Section 2.1, "Starting the Configuration Wizard in Graphical Mode."

   The **Welcome** screen is displayed.

2. Select **Create a new WebLogic domain** and click **Next**.

   The **Select Domain Source** screen is displayed.

3. Select one of the following options:

   - **Generate a domain configured automatically to support the following products**:

     The **WebLogic Server** option is selected, by default.

     Select the check boxes corresponding to the other products to include in the domain.

   - **Base this domain on an existing template**

     Select this option to create a domain by using an existing domain template. For information about domain templates, see *WebLogic Server Domain Template Reference*.

     Enter the full path to the template in the **Template location** field, or click **Browse** to navigate to the directory containing the required template.

   Click **Next**.

   The **Specify Domain Name and Location** screen is displayed.

4. Enter the name of the domain and specify the domain location.

   - Domain names must not start with a number. This restriction is intended to prevent potential conflicts with internally-generated JDBC store table names, which must begin with a letter.
   - The domain directory can be located anywhere in the system. By default, it resides in `MW_HOME\user_projects\domains\domain`, where `MW_HOME` is the directory that contains the product installation, and `domain` is the name of the domain directory defined by the selected template.

     The Configuration Wizard stores the `config.xml` file and all other generated components in the domain directory that you specify.

   Click **Next**.

The **Configure Administrator User Name and Password** screen is displayed.

5. Configure the username and password for the administrator. The username is used to boot the administration server and connect to it.

   - The user name **must not** contain commas, tabs, or any of the following characters: `<>#|&?(){}`.
   - The password is case sensitive, and **must** contain at least eight alphanumeric characters, with at least one numeral or one of the following characters: `!"#$%&'()*+,-./:;<=>?@[\]^_`{|}~`

   Click **Next**.

   The **Configure Server Start Mode and JDK** screen is displayed.

6. Select the WebLogic domain startup mode.

   - In the **development** mode, the configuration of security is relatively relaxed, allowing you to autodeploy applications.
   - In the **production** mode, the configuration of security is stringent, requiring a user name and password to deploy applications. Before putting a domain into production, familiarize yourself with the securing the production environment. For more information, see *Securing a Production Environment*.

   For information about changing the run-time mode after you have created a domain, see the *WebLogic Server Administration Console Online Help*.

   Table 3–1 provides information to help you select a startup mode that suits your requirements.

Table 3–1 Differences Between Development Mode and Production Mode

Function	Development mode	Production mode
SSL	You can use the demonstration digital certificates and the demonstration keystores provided by the WebLogic Server security services. With these certificates, you can design your application to work within environments secured by SSL. For more information, see *Securing Oracle WebLogic Server*.	You must not use the demonstration digital certificates and the demonstration keystores. If you do so, a warning message is displayed.
Deploying applications	WebLogic Server instances can deploy and update applications that reside in the `domain_name/autodeploy` directory automatically. It is recommended that this method be used only in a single-server development environment. For more information, see *Deploying Applications to WebLogic Server*.	The auto-deployment feature is disabled; so, you must use the WebLogic Server administration console, the `weblogic.Deployer` tool, or the WebLogic Scripting Tool.
Log file rotation	By default, when you start the WebLogic Server instance, the server automatically renames (rotates) its local server log file as `SERVER-NAME.log.n`. For the remainder of the server session, messages accumulate in the log file until the file grows to a size of 500 kilobytes. For more information, see the "Rotate Log Files" topic in the *WebLogic Server Administration Console Online Help*.	The server rotates the local log file after the size of the file reaches 5000 kilobytes. When the server is configured for production mode, by default, all versions of the log files are kept. Administrators may want to customize the number of log files retained.
JDBC system resource	The default capacity is 15 connections.	The default capacity is 25 connections.

7. Select the JDK.

   In the **JDK Selection** pane, select the JDK for the startup mode that you selected in the **WebLogic Domain Startup Mode** pane.

   The Configuration Wizard presents a list of the JDKs included in the installer. You can choose one of these JDKs or another JDK that you have installed on your system.

   > **Note:** If you select a JDK that is included in the installer, the Configuration Wizard creates server startup scripts to invoke the JDK. If you select a JDK that is not supplied by JDK, the Configuration Wizard does not configure the startup scripts; you must change the startup scripts manually.

   Select only those JDKs that are supported on the platform you are using. For a list of the JDKs that are supported for a specific platform, see Oracle Fusion Middleware Supported System Configurations at `http://www.oracle.com/technology/software/products/ias/files/fusion_certification.html`. The default selection reflects the JDK that best meets the requirements of your environment, based on the platform on which you are installing the domain.

8. After selecting the JDK, click **Next**.

   - If the domain includes components that require JDBC data sources, the **Configure JDBC Data Source** screen is displayed.

     For more information, see Section 5.1, "Configure JDBC Data Sources."

   - If the domain includes components (such as, SOA) that require the use of the Oracle Fusion Middleware Repository Creation Utility (RCU) to load database schemas, the **Configure JDBC Component Schema** screen is displayed.

     For more information, see Section 5.5, "Configure JDBC Component Schema."

   If the domain that you are creating does not include any component that requires JDBC data sources or RCU-loaded database schemas, the **Select Advanced Configuration** screen is displayed.

9. While creating the domain, you can perform certain advanced configuration tasks, such as modifying settings for the administration server; configuring the distribution of your domain across servers, clusters, and machines; configuring JMS file stores; and configuring RDBMS security store settings.

   In the **Select Advanced Configuration** screen, select the appropriate check boxes depending on the configuration tasks to perform, and click **Next**. For more information about advanced configuration options, see Chapter 4, "Customizing the Domain Environment."

   If you click **Next** without selecting any advanced configuration option, the **Configuration Summary** screen is displayed.

10. Review the configuration settings of your domain before the Configuration Wizard creates the domain.

    Select an item in the **Domain Summary** pane on the left and review the associated details in the **Details** pane on the right. You can make limited adjustments by clicking **Previous** to return to a previous screen.

Note: You can limit the type of information displayed in the Domain Summary pane by selecting a filter from the Summary View drop-down list.

11. After reviewing the domain settings, click **Create**.

    The **Creating Domain** screen is displayed. It displays messages indicating the progress of the domain creation process.

    When the process is complete, the new domain is ready for use.

12. To start the administration server immediately, select the **Start Admin Server** check box. This option is available only for Windows systems.

13. Click **Done**.

4

Customizing the Domain Environment

While creating a domain, you can (optionally) perform certain advanced configuration tasks such as the following:

- Modifying settings for the administration server
- Configuring the distribution of your domain across servers, clusters, and machines
- Assigning deployments and servers to clusters or servers
- Configuring JMS file store settings
- Configuring the RDBMS security store

To perform the advanced configuration tasks, you must select the appropriate check boxes in the **Select Advanced Configuration** screen.

The Configuration Wizard guides you through the following configuration tasks.

4.1 Configuring the Administration Server

In every domain, one server must be designated as the administration server: the central point from which the whole domain is managed.

You can access the administrator server by using the URL `protocol://listen-address:listen-port`. The `protocol` can be any of the following: `t3`, `t3s`, `http`, `https`.

Table 4–1 describes the fields in the **Configure the Administration Server** screen. Specify the appropriate values, and then click **Next** in the wizard to proceed.

Note: Fields marked with an asterisk (*) are mandatory fields.

Table 4–1 Configuring the Administration Server

Field	Description
Name*	Enter a valid server name: a string of characters. Each server instance in the production environment must have a unique name, regardless of the domain or cluster in which it resides, and regardless of whether it is an administration server or a managed server. In addition, the name of the administration server must be unique among all component names within the domain. **Note**: This value is specified for identification purposes only. It is not used as part of the URL for applications that are deployed on the server. The server name is displayed in the WebLogic Server administration console. In addition, if you use WebLogic Server command-line utilities or APIs, you must specify this name to identify the server.
Listen address	From the drop-down list, select a value for the listen address. If you select **localhost** as the listen address for a server instance, remote processes cannot connect to that server instance. Only processes on the machine that hosts the server instance can connect to the server instance. If the server instance must be accessible as localhost (for example, if you create administrative scripts that connect to localhost), and it must also be accessible by remote processes, select **All Local Addresses**. The server instance determines the address of the machine and listens on it. For more information, see Section 4.1.1, "Specifying the Listen Address."
Listen port	Enter a valid value for the listen port to be used for regular, nonsecure requests (through protocols such as HTTP and T3). The default value is 7001 for the administration server. If you leave this field blank, the default value is used. The valid listen port range is from 1 to 65535. For more information, see Section 4.1.2, "Specifying the Listen Port."
SSL enabled	Select this check box to enable the SSL listen port. By default, SSL is disabled for all new servers.
SSL listen port	This field is enabled only if the **SSL enabled** check box is selected. Enter a valid value to be used for secure requests (through protocols such as HTTPS and T3S). The default value is 7002. If you leave this field blank, the default value is used. The valid listen port range is from 1 to 65535. **Note**: By default, a server instance uses demonstration certificates to authenticate requests from a secure port. In a production environment, you must configure SSL to use certificates from a certificate authority. For more information, see Configuring SSL. For more information, see Section 4.1.2, "Specifying the Listen Port."

4.1.1 Specifying the Listen Address

Table 4–2 provides guidelines for specifying the listen address for a server.

Table 4–2 Specifying Listen Address

Listen Address	Behavior
All Local Addresses or a DNS name	On multi-homed Windows machines, a server instance binds to all available IP addresses.
An IP address or a DNS name	■ To connect to the server instance, processes can specify either the IP address or the corresponding DNS name. ■ Processes that specify localhost fail to connect. ■ You must update existing processes that use localhost to connect to the server instance. ■ For connections that specify the IP address for the listen address and a secured port for the listen port, host name verification must be disabled. **Note**: To resolve a DNS name to an IP address, WebLogic Server must be able to contact an appropriate DNS server or obtain the IP address mapping locally. Therefore, if you specify a DNS name for the listen address, you must either leave a port open long enough for the WebLogic Server instance to connect to a DNS server and cache its mapping or you must specify the IP address mapping in a local file. If you specify an IP address for the listen address and then a client request specifies a DNS name, WebLogic Server attempts to resolve the DNS name, but if it cannot access DNS name mapping, the request fails.
localhost	■ Processes must specify localhost to connect to the server instance. ■ Only processes that reside on the machine that hosts the server instance (local processes) can connect to the server instance.

4.1.2 Specifying the Listen Port

Note the following guidelines when specifying the listen ports and secure listen port:

- Although you can specify any valid port number, if you specify port 80, you can omit the port number from the HTTP request used to access resources over HTTP. For example, if you define port 80 as the listen port, you can use the URL: `http://hostname/myfile.html` instead of `http://hostname:portnumber/myfile.html`.
- On some operating systems, port 80 can be accessed only by processes run under a privileged user or group ID. In this case, you can assign the server instance to a UNIX machine on which a post-bind UID or GID is defined.

4.2 Configuring Managed Servers

In production environments, enterprise applications are hosted, typically, on one or more managed servers, in addition to the administration server.

You can add and delete managed servers in the **Configure Managed Servers** screen of the Configuration Wizard.

> **Note:** You can create managed servers on remote machines by using the `pack` and `unpack` commands. For more information, see the *Creating Templates and Domains Using the Pack and Unpack Commands*.

1. Review the current managed server configurations. Default values may vary based on the domain source you selected earlier.

Note: The wizard provides two views: a concise tabular view of all the managed servers and an individual view of each managed server with each server represented by a tab. To toggle the view mode between table and tab views, click **Switch Display**.

2. Add or delete managed servers, or change the settings for existing managed servers, as required for your domain.

   The screens in this screen are identical to those in the **Configure the Administration Server** screen. For more information, see Section 4.1, "Configuring the Administration Server."

3. After configuring the managed servers, click **Next** in the wizard to proceed.

4.3 Configuring Clusters

A cluster is a group of WebLogic Server instances that work together to provide scalability and high-availability for applications. By creating clusters, you can group managed servers such that they operate as a single unit for hosting applications and resources.

You can add, configure, and delete clusters in the **Configure Clusters** screen of the Configuration Wizard. This screen is displayed only if the domain contains at least one managed server.

1. Review the current cluster configuration. Default values may vary, based on the domain source you selected earlier.

   Note: The wizard provides two views: a concise tabular view of all the clusters, and an individual view of each cluster with each cluster represented by a tab. To toggle the view mode between table and tab views, click **Switch Display**.

2. Add or delete clusters, or change the settings for existing clusters, as required for your domain.

   Note: Fields marked with an asterisk (*) are mandatory fields.

Table 4–3 Configuring Clusters

Field	Action
Name*	Enter a valid name for the cluster: a string of characters. The name of the cluster must be unique among all component names within the domain. The default value in this field is `new_Cluster_n`, where `n` is a numeric value that is used to differentiate among all the default cluster names; the value of n for the first cluster is 1. The value is incremented by 1 for each cluster that you add.
Cluster messaging mode	Select the cluster message mode: unicast or multicast.

Table 4–3 (Cont.) Configuring Clusters

Field	Action
Multicast address	If you selected **multicast** as the cluster message mode, enter the multicast address for the cluster. This address is used by cluster members to communicate with each other. The default value is 239.192.0.0. The valid multicast address range is 224.0.0.1 to 239.255.255.255.
Multicast port	f you selected **multicast** as the cluster message mode, enter the multicast port for the cluster. The multicast port is used by cluster members to communicate with each other. The default value is 7001. Valid values for multicast ports are from 1 to 65535.
Cluster address	Enter the addresses to identify the managed servers in the cluster. A cluster address can be one of the following: ■ Comma-separated list of IP addresses or DNS names and ports (for example: `dns_name:port,dns_name:port`) ■ DNS name that maps to multiple IP addresses ■ `localhost`, DNS name, or IP address if the listen address of all managed servers is listening to the same address with unique port numbers The cluster address is used in entity and stateless EJBs to construct the host name portion of URLs. If the cluster address is not set, EJB handles may not work properly.

3. After configuring the clusters, click **Next** in the wizard to proceed.

For more information, see *Using WebLogic Server Clusters*.

4.4 Assigning Managed Servers to Clusters

You can assign the available managed servers to clusters within the domain in the **Assign Servers to Clusters** screen.

This screen is displayed only if you have defined at least one cluster.

1. In the **Cluster** pane, select the cluster to which you want to assign a managed server.
2. Assign managed servers to the selected cluster in one of the following ways:
   - Double-click the name of the managed server in the **Server** pane.
   - Select the managed server and click the right arrow.
   - Shift+click to select multiple managed servers; then, click the right arrow.

   The name of the managed server is removed from the **Server** pane and added below the name of the target cluster in the **Cluster** pane.

   Note: Only managed servers are listed in the **Server** pane. The administration server is not listed because it cannot be assigned to a cluster.

3. Repeat steps 1 and 2 for each managed server to assign to a cluster.
4. Review the cluster assignments.

If necessary, you can remove a managed server from a cluster in one of the following ways:

- Double-click the name of the managed server in the **Cluster** pane.
- Select the managed server and click the left arrow.

The name of the managed server is removed from the **Cluster** pane and restored to the **Server** pane.

5. Click **Next** in the wizard to proceed.

4.5 Creating HTTP Proxy Applications

An HTTP proxy application acts as an intermediary for HTTP requests.

In the **Create HTTP Proxy Applications** screen of the Configuration Wizard, you can create an HTTP proxy application for each cluster, and specify the managed server on which the proxy application must be deployed.

This screen is displayed only if both of the following statements are true:

- At least one managed server is assigned to a cluster.
- At least one managed server is not assigned to any cluster.

To create HTTP proxy applications:

1. Select the **Create HTTP Proxy** check box against the cluster for which you want to create the HTTP proxy application.

   A list of the managed servers that are not assigned to any cluster is displayed in the **Proxy Server** list.

2. From the **Proxy Server** list, select a managed server on which the proxy application must be deployed.

   A proxy application named `OracleProxy4_<clustername>_<servername>` is created and deployed in the managed server.

3. Repeat steps 1 through 3 for each cluster for which you want to create HTTP proxy applications.
4. Click **Next** in the wizard to proceed.

4.6 Configuring Machines

In a domain, the machine definitions identify physical units of hardware and are used to associate computers with the WebLogic Server instances that they host.

You might want to create machine definitions in situations such as the following:

- The administration server uses the machine definition, with the node manager application, to start remote servers.
- WebLogic Server uses configured machine names when determining the server in a cluster that is best able to handle certain tasks, such as HTTP session replication. WebLogic Server then delegates those tasks to the identified server.

Note: You must configure machines for each product installation that runs a node manager process. The machine configuration must include values for the listen address and port number parameters.

You can create machine definitions in the **Configure Machines** screen of the Configuration Wizard.

1. Select the **Machine** tab (for Windows) and **Unix Machine** tab (for UNIX).
2. Review the current list of configurations, and add or change settings as required for your domain.
   - To add a machine, click **Add**.
   - To delete a machine, select the machine in the list and click **Delete**.

   Table 4–4 describes the configuration settings that you can define. Default values may vary based on the domain source that you selected earlier.

Note: Fields marked with an asterisk (*) are mandatory fields.

Table 4–4 Configure Windows Machine

Field	Description
Name*	Enter a valid machine name: a string of characters. The machine name is used to identify the machine within the WebLogic Server domain; it is not required to match the network name for the machine. The name must be unique among all component names within the domain. The default value in this field is `new_[Unix]Machine_n`, where `n` is a numeric value that is used to differentiate among all default machine names; the value of `n` for the first machine is 1. The value is incremented by 1 for each machine that you add.
Node manager listen address	Select a value from the drop-down list for the listen address used by node manager to listen for connection requests. By default, the IP addresses defined for the local system and localhost are shown in the drop-down list. The default value is localhost. If you specify an IP address for a machine that hosts the administration server and you need to access the WebLogic Server node manager, you must disable host name verification. For more information, see "Using Host Name Verification" in *Securing Oracle WebLogic Server*.
Node manager listen port	Enter a valid value for the listen port used by node manager to listen for connection requests. The valid node manager listen port range is from 1 to 65535. The default value is 5556.
Post bind GID enabled	This field is displayed only in the **Unix Machine** tab. Select this check box to enable a server running on this machine to bind to a UNIX group ID (GID) after it finishes all privileged startup actions. By default, this check box is not selected.
Post bind GID	This field is displayed only in the **Unix Machine** tab. Enter a valid UNIX group ID (GID) that a server running on this machine will run under after it finishes all privileged startup actions. Otherwise, the server continues to run under the group from which it was started. For this setting to take effect, you must select the **Post bind GID enabled** check box.
Post bind UID enabled	This field is displayed only in the **Unix Machine** tab. Select this check box to enable a server running on this machine to bind to a UNIX user ID (UID) after it finishes all privileged startup actions. By default, this check box is not selected.

Table 4–4 (Cont.) Configure Windows Machine

Field	Description
Post bind UID	This field is displayed only in the **Unix Machine** tab. Enter a valid UNIX user ID (UID) that a server running on this machine will run under after it finishes all privileged startup actions. Otherwise, the server continues to run under the account from which it was started. For this setting to take effect, you must select the **Post bind UID enabled** check box.

3. After updating the settings, click **Next**.

The **Assign Servers to Machines** screen is displayed.

4.7 Assigning Servers to Machines

After configuring servers and defining machines, you can assign WebLogic Server instances to machines in the **Assign Servers to Machines** screen.

This screen is displayed only if you have defined at least one machine.

1. In the **Machine** pane, select the Windows or UNIX machine to which you want to assign a WebLogic Server instance.
2. Assign WebLogic Server instances to the selected machine in one of the following ways:
   - Double-click the WebLogic Server instance in the **Server** pane.
   - Select the appropriate WebLogic Server instance in the **Server** pane and click the right arrow.
   - Shift+click to select multiple servers in the **Server** pane; then, click the right arrow.

   The name of the WebLogic Server instance is removed from the **Server** pane and added, below the name of the target machine, in the **Machine** pane.
3. Repeat steps 1 and 2 for each WebLogic Server instance to assign to a machine.
4. Review the machine assignments.

   If necessary, you can remove a WebLogic Server instance from a machine in one of the following ways:
   - Double-click the name of the appropriate WebLogic Server instance in the **Machine** pane.
   - Select the appropriate WebLogic Server instance in the **Machine** pane and click the left arrow.

   The name of the WebLogic Server instance is removed from the **Machine** pane and restored to the **Server** pane.
5. Click **Next**.

4.8 Target Deployments to Clusters or Servers

The **Target Deployments to Clusters or Servers** screen of the Configuration Wizard is displayed only if the selected template contains J2EE applications or libraries.

In this screen you can target applications and libraries for deployment on servers or clusters.

Applications and libraries associated with the product for which you are configuring the domain (for example, SOA) are targeted automatically, to the managed server created for that product or to the cluster to which that managed server is assigned. In this screen, you can target applications and libraries to additional servers and clusters.

Note: When you extend a domain, if you delete a managed server or cluster to which applications are currently targeted, the Configuration Wizard automatically retargets those applications as follows:

- If the applications were originally targeted solely to the managed server or cluster that you are now deleting (that is, after you delete the managed server or cluster, the applications would become untargeted in the modified domain), then the Configuration Wizard automatically retargets the applications to all **eligible** targets.

  An eligible target is any cluster or managed server that is not defined in the configuration groups specification (`config-groups.xml` file) of an included template. Servers or clusters that are specified in `config-groups.xml` are essentially **owned** by the template and, therefore, are not eligible for automatic targeting.

- If the applications were originally targeted to multiple targets (including managed servers, clusters, and the administration server), and one of the targeted managed servers or clusters is deleted, then, in the extended domain, the Configuration Wizard leaves the remaining target associations intact and does not attempt to retarget the applications.

---

1. In the **Target** pane, select the cluster or server on which you want to deploy applications or libraries.

   The name of the selected target is displayed as the title of the right pane.

2. In the right pane, select the check boxes corresponding to the applications or libraries to deploy on the selected target.

   The applications and libraries displayed here vary, depending on the products that you selected in the **Select Domain Source** screen, earlier in the wizard.

   Note: When you select a managed server in the **Target** pane, some of the check boxes in the right pane might be disabled, indicating applications and libraries that are already targeted at the cluster that contains the selected managed server.

   After you select applications and libraries, the names of the targeted clusters and servers are displayed in the **Target** column in the right pane.

3. Repeat steps 1 and 2 for the other clusters and servers, as required.

4. After making the required selections, click the **Next** button in the wizard.

4.9 Target Services to Clusters or Servers

The **Target Services to Clusters or Servers** screen of the Configuration Wizard is displayed only if the selected template contains J2EE services.

Services that are associated with the product for which you are configuring the domain (for example, SOA) are targeted automatically, to the managed server created for that product or to the cluster to which that managed server is assigned. In this screen, you can target services to additional servers and clusters.

Note: When you extend a domain, if you delete a managed server or cluster to which services are currently targeted, the Configuration Wizard automatically retargets those services as follows:

- If the services were originally targeted solely to the managed server or cluster that you are now deleting (that is, after you delete the managed server or cluster, the services would become untargeted in the modified domain), then the Configuration Wizard automatically retargets the services to all **eligible** targets.

  An eligible target is any cluster or managed server that is not defined in the configuration groups specification (`config-groups.xml` file) of an included template. Servers or clusters that are specified in `config-groups.xml` are essentially **owned** by the template and, therefore, are not eligible for automatic targeting.

- If the services were originally targeted to multiple targets (including managed servers, clusters, and the administration server), and one of the targeted managed servers or clusters is deleted, then, in the extended domain, the Configuration Wizard leaves the remaining target associations intact and does not attempt to retarget the services.

---

1. In the **Target** pane, select the cluster or server on which you want to deploy services.

   The name of the selected target is displayed as the title of the right pane.

2. In the right pane, select the check boxes corresponding to the services to deploy on the selected target.

   The services displayed here vary, depending on the products that you selected in the **Select Domain Source** screen earlier in the wizard.

   Note: When you select a managed server in the **Target** pane, some of the check boxes in the right pane might be disabled, indicating services that are already targeted at the cluster that contains the selected managed server.

   After you select services, the names of the targeted clusters and servers are displayed in the **Target** column in the right pane.

3. Repeat steps 1 and 2 for the other clusters and servers, as required.

4. After making the required selections, click the **Next** button in the wizard.

4.10 Configure JMS File Stores

A JMS file store is a disk-based file in which persistent messages can be saved.

You can modify the JMS file stores that are configured in your domain, in the **Configure JMS File Stores** screen, which is displayed only if the selected template contains JMS.

1. Review the current list of JMS file stores. Default values may vary based on the domain source that you selected earlier.

   Note: The wizard provides two display modes: a concise tabular view of all the defined components, and an individual view in which each component is represented by a tab. To toggle the view mode between table and tab views, click **Switch Display**.

2. Modify the settings, as required for your domain.

   Table 4–5 describes the fields of the **Configure JMS File Stores** screen.

   Note: Fields marked with an asterisk (*) are mandatory fields.

Table 4–5 Configure JMS File Stores

Field	Description
Name*	Enter a name for the JMS file store. The name must be a string of characters. The name of the JMS file store must be unique among all component names within the domain.
Directory	Enter the path of the directory (in your system) in which the JMS file store resides.
Synchronous write policy	From the drop-down list, select one of the following synchronous write policies to determine how the file store writes data to the disk: ■ **Cache-Flush**: Transactions cannot be completed until all their write operations have been flushed to the disk. ■ **Direct-Write**: Write operations are performed directly to the disk. This policy is supported on Solaris and Windows. If this policy is active on an unsupported platform, the file store switches automatically to the cache-flush policy. ■ **Disabled**: Transactions are complete as soon as the writes are cached in memory. When this policy is active, completion of transactions does not depend on waiting for writes to reach the disk. This setting affects performance, scalability, and reliability. **Note:** The use of the direct-write policy is reliable in Solaris systems, but Windows systems may leave transaction data in the on-disk cache without writing it to disk immediately. This is considered unreliable, because a power failure can cause loss of on-disk cache data, possibly resulting in lost or duplicate messages. For reliable writes using the direct-write policy on Windows, either disable all write caching for the disk (enabled by default), or use a disk with a battery-backed cache. Some file systems, however, do not allow this value to be changed (for example, a RAID system that has a reliable cache). **Note:** If the JMS file store is used exclusively for paging nonpersistent messages to the disk, the synchronous write policy is ignored.

3. After updating the settings, click **Next**.

The **Configure RDBMS Security Store Database** screen is displayed.

4.11 Configuring the RDBMS Security Store Database

You can define RDBMS security store settings in the **Configure RDBMS Security Store Database** screen of the Configuration Wizard.

Table 4–6 describes the fields in the **Configure RDBMS Security Store Database** screen.

> **Note:** Fields marked with an asterisk (*) are mandatory fields.

Table 4–6 Configure RDBMS Security Store Database

Field	Description
*Database Type	From the drop-down list, select the type of database to use as the RDBMS security store.
*Driver	Select the driver to use for the database. The list of available drivers varies, depending on the database type that you select.
*Class Name	No action is required. The class name is displayed automatically, based on the driver that you select.
*DBMS SID	Enter the SID of the database.
*DBMS Host	Enter the name of the machine that hosts the database.
*DBMS Port	Enter the port to be used to connect to the server. The default port number that is associated with the selected database type is displayed automatically.
*URL	No action is required. The URL is displayed automatically based on the driver that you select.
*User Name	Enter the login name for connecting to the database.
*User Password	Enter the password for accessing the database. The password rules vary, depending on the database. The value is encrypted.
*Confirm User Password	Re-enter the password.
*Known Properties	No action is required. The known properties of the database are displayed automatically based on the driver that you select.
Additional Properties	Enter additional properties, if any, to be passed to the driver.

After specifying the RDBMS security store settings, click **Next**.

> **Note:** You can test the connection to the database by clicking **Test Connection**. Before starting the server, you must load the necessary SQL scripts for the RDBMS security store. If you use an RDBMS security store in a clustered domain, it is recommended that you use it with JMS configuration (JNDI name and JMS topic). For more information, see the *WebLogic Server Administration Console Online Help*.

The **Configuration Summary** screen is displayed.

4.12 Review the Domain Settings and Create the Domain

In the **Configuration Summary** screen, you review the detailed configuration settings of your domain before the Configuration Wizard creates it.

1. Select an item in the **Domain Summary** pane on the left and review the associated details in the **Details** pane on the right. You can make limited adjustments by clicking **Previous** to return to a previous screen.

   > **Note:** You can limit the type of information displayed in the Domain Summary pane by selecting a filter from the **Summary View** drop-down list.

2. After reviewing the domain settings, click **Create**.

   The **Creating Domain** screen is displayed. It displays messages indicating the progress of the domain creation process.

   When the process is complete, the new domain is ready for use.

3. To start the server immediately, select the **Start Admin Server** check box and click **Done**. This option is available only in Windows systems.
4. Click **Done**.

5

Customizing JDBC Data Sources and Component Schema

When you are creating or extending a domain using the Configuration Wizard, you can change the JDBC data source and JDBC component schema settings if they are defined in the domain or template that you selected as the source for domain that you are creating or extending.

The following topics describe the steps required to change the JDBC settings in your domain:

5.1 Configure JDBC Data Sources

A JDBC data source contains a pool of database connections that are created when the data source instance is created—when it is deployed or targeted, or at server startup. Applications look up a data source on the JNDI tree, and then request a connection. When the applications no longer need the connections, they return the connections to the connection pool in the data source.

In the **Configure JDBC Data Sources** screen, you can configure the JDBC data sources defined in your domain source.

The JDBC data sources associated with the products for which you are creating the domain are listed in the lower half of the screen.

> **Note:** To configure one or more of the data sources as RAC multi data sources, select the check box adjacent to the name of the required schema, and select the **Configure selected data sources as RAC multi data sources in the next panel** check box; then, click the **Next** button.

1. Select the data source for which you want to specify settings by selecting the check box adjacent to the data source name.

Note: When you select multiple data sources, the text "`Varies among data sources`" might be displayed in certain fields, indicating that the current values of those fields are different across the selected data sources. If you go ahead and change the values in such fields, the new values are applied uniformly across the selected data sources.

2. Review the current configuration settings and modify them as required. Table 5–1 describes the fields of the **Configure JDBC Data Sources** screen.

Note: The default values of data source parameters such as vendor, driver, host name, and port number depend on the values that are specified in the application templates.

Table 5–1 Configure JDBC Data Sources

Field	Description
Vendor	Select the database vendor.
Driver	Select the JDBC driver to use to connect to the database. The list includes common JDBC drivers for the selected database vendor.
Username	Enter the account login name for connecting to the database.
Password	Enter the login password for the specified username.
DBMS/Service	Enter a DBMS SID or service name. The value that you must enter depends on the driver that you select. ▪ If the name of the Oracle driver that you selected contains the words "for Instance connections," you must enter the SID. ▪ If the name of the Oracle driver contains the words "for Service connections," you must enter the service name. ▪ For information about configuring DataDirect's driver, see the DataDirect documentation.
Host Name	Enter the name of the server hosting the database.
Port	Enter the port number to be used to connect to the server.

The values that you specify are displayed in the appropriate columns in the data source list, for the selected data source.

3. After updating the settings, do one of the following:
   - To test the data source connections to the specified databases, click **Test Connections**. The **Test Data Source Connections** screen is displayed. For more information, see Section 5.3, "Test JDBC Data Sources."
   - If you do not want to test the connections to the database, click **Next**.

5.2 Configure RAC Multi Data Source

In this screen of the Configuration Wizard, you configure the data sources (that are included in the domain) as RAC data sources.

The data sources that you opted to configure as RAC data sources in the previous screen of the wizard (**Configure JDBC Data Sources**) are listed in the lower half of the screen.

1. Select the data source for which you want to specify data source settings by selecting the check box adjacent to the data source name.

   Note: When you select multiple data sources, the text "`Varies among data sources`" might be displayed in certain fields, indicating that the current values of those fields are different across the selected data sources. If you go ahead and change the values in such fields, the new values are applied uniformly across the selected data sources.

2. Specify the settings for the RAC data sources.

   Table 5–2 describes the fields in this screen.

Table 5–2 Configure JDBC Data Sources

Field	Description
Driver	Select the JDBC driver to use to connect to the database.
Service Name	Enter a database Oracle RAC service name.
Username	Enter the account login name for connecting to the database.
Password	Enter the login password for the specified username.
Host Name	Enter the name of the server hosting the RAC database instances.
Instance Name	Enter the name of each Oracle database instance.
Port	Enter the port numbers to be used to connect to the server that hosts the database.

   Note: You must specify the host name, instance name, and port number of at least one database instance.

   To add a database instance, click **Add**, and then specify the host name, instance name, and port.

   The values that you specify are displayed in the appropriate columns in the data source list, for the selected data source.

3. After making the required selections and entries, click the **Next** button in the wizard to test the connections to the data sources.

5.3 Test JDBC Data Sources

In the **Test JDBC Data Sources** screen, you can test the configurations that you specified for the data sources in the previous screen.

1. Select the check boxes adjacent to the names of the data sources to test.
2. Click **Test Connections** for the connection to test.

The wizard tests the configuration for each data source, by attempting to connect to a URL that is constructed by using the driver, host, port, and other information that you specified while configuring the data source.

The result of the test is indicated in the **Status** column. Details are displayed in the **Connection Result Log** pane.

3. After testing the connections, click the **Next** button in the wizard.

5.4 Run Database Scripts

A domain template might contain a set of SQL files organized by database type. If the domain template contains SQL files, you can run them while creating the domain, in the **Run Database Scripts** window. Database content for each of the data sources defined in your domain is set up by using pre-existing SQL or database loading files.

Note: No databases are defined in the `wls.jar` template; so, if you selected the WebLogic Server template as the basis for the domain, the **Configure JDBC Data Sources** window and the **Run Database Scripts** screens are not displayed.

---

1. In the **Available JDBC Data Sources** pane, select the data source for which you want to run the scripts. The scripts that can be executed are displayed in the **Available SQL Files and Database Loading Options**.
2. Select the database version from the **DB Version** drop-down menu.
3. Click **Run Scripts**.

   All the scripts displayed in the **Available SQL Files and Database Loading Options** pane for the selected data source are executed, and the results are displayed in the **Results** pane. To capture test output in a log file, select the **Log File** check box and specify the location of the log file.
4. Repeat steps 1 through 3 for each data source for which you want to execute SQL scripts.
5. Click **Next**.

5.5 Configure JDBC Component Schema

For some Fusion Middleware components (for example, SOA and WebCenter Spaces), JDBC data sources might be defined as part of the component's database schema, which are loaded (during installation) by using the Repository Creation Utility (RCU).

When you create a domain for such components by using the Configuration Wizard, you can configure the JDBC component schema settings: database driver, schema owner, password, and so on.

The JDBC component schemas associated with the products for which you are creating the domain (example, SOA) are listed in the lower half of the screen.

Note: To configure one or more of the schemas as RAC multi data sources, select the check boxes adjacent to the name of the required schemas, and select the **Configure selected component schemas as RAC multi data source schemas** in the next panel check box; then, click the **Next** button.

---

1. Select the schemas for which you want to specify data source settings, by selecting the check box adjacent to the schema names.

> **Note:** When you select multiple component schemas, the text "`Varies among component schemas`" might be displayed in certain fields, indicating that the current values of those fields are different across the selected component schemas. If you go ahead and change the values in such fields, the new values are applied uniformly across the selected component schemas.

2. Review the current configuration settings and modify them as required.
3. The following table describes the fields in this screen.

> **Note:** The default values of component schema parameters such as vendor, driver, host name, and port number depend on the values that are specified in the application templates.

Table 5–3 Configure JDBC Data Sources

Field	Description
Vendor	Select the database vendor.
Driver	Select the JDBC driver to use to connect to the database. The list includes common JDBC drivers for the selected database vendor.
Schema Owner	Enter the account login name for connecting to the database.
Schema Password	Enter the login password for the specified schema owner.
DBMS/Service	Enter a database DBMS name, or service name if you selected a service type driver.
Host Name	Enter the name of the server hosting the database.
Port	Enter the port number to be used to connect to the server that hosts the database.

The values that you specify are displayed in the appropriate columns in the schema list, for the selected schemas.

4. After making the required selections and entries, click the **Next** button in the wizard to test the component schemas.

5.6 Configure RAC Multi Data Source Component Schema

In this screen of the Configuration Wizard, you configure the component schemas (that are included in the domain) as RAC multi data sources.

The component schemas that you opted to configure as RAC multi data sources in the previous screen of the wizard (**Configure JDBC Component Schema**) are listed in the lower half of the screen.

1. In the schema list in the lower half of the screen, select the schemas to configure as RAC multi data sources, by selecting the check box adjacent to the schema names.

Note: When you select multiple data source schemas, the text "`Varies among component schemas`" might be displayed in certain fields, indicating that the current values of those fields are different across the selected schemas. If you go ahead and change the values in such fields, the new values are applied uniformly across the selected schemas.

2. Specify the settings for the data sources.

   Table 5–2 describes the fields in this screen.

Table 5–4 Configure JDBC Data Sources

Field	Description
Driver	Select the JDBC driver to use to connect to the database.
Service Name	Enter a database Oracle RAC service name.
Username	Enter the account login name for connecting to the database.
Password	Enter the login password for the specified username.
Host Name	Enter the name of the server hosting the RAC database instances.
Instance Name	Enter the name of each Oracle database instance.
Port	Enter the port numbers to be used to connect to the server that hosts the database.

Note: You must specify the host name, instance name, and port number of at least one database instance.

   To add a database instance, click **Add**, and then specify the host name, instance name, and port.

   The values that you specify are displayed in the appropriate columns in the schema list, for the selected schemas.

3. After making the required selections and entries, click the **Next** button in the wizard to test the component schemas.

5.7 Test Component Schema

In the **Test Component Schema** screen, you can test the configurations that you specified for the data sources in the previous screen.

1. Select the check boxes adjacent to the names of the schemas to test.
2. Click **Test Connections** for the connection to test.

   The wizard tests the configuration for each schema, by attempting to connect to a URL that is constructed by using the driver, host, port, and other information that you specified while configuring the schema.

   The result of the test is indicated in the **Status** column. Details are displayed in the **Connection Result Log** pane.

3. After testing the connections, click the **Next** button in the wizard.

5.8 Configure JMS File Stores

A JMS file store is a disk-based file in which persistent messages can be saved.

You can modify the JMS file stores that are configured in your domain, in the **Configure JMS File Stores** screen, which is displayed when you click **Next** in the **Run Database Scripts** screen. This step is optional.

1. Review the current list of JMS file stores. Default values may vary based on the domain source that you selected earlier.

   Note: The wizard provides two display modes: a concise tabular view of all the defined components, and an individual view, in which each component is represented by a tab, and you view a particular component by selecting the corresponding tab. To toggle the display mode between table and tab formats, click Switch Display.

2. Modify the settings, as required for your domain.

   Table 5–1 describes the fields of the **Configure JMS File Stores** screen.

   Note: Fields marked with an asterisk (*) are mandatory fields.

Table 5–5 Configure JMS File Stores

Field	Description
Name*	Enter a name for the JMS file store. The name must be a string of characters and can include spaces. The name of the JMS file store must be unique among all component names within the domain.
Directory	Enter the path of the directory (in your system) in which the JMS file store resides.

Table 5–5 (Cont.) Configure JMS File Stores

Field	Description
Synchronous write policy	From the drop-down list, select one of the following synchronous write policies to determine how the file store writes data to the disk: ■ **Cache-Flush**: Transactions cannot be completed until all their write operations have been flushed to the disk. ■ **Direct-Write**: Write operations are performed directly to the disk. This policy is supported on Solaris and Windows. If this policy is active on an unsupported platform, the file store switches automatically to the cache-flush policy. ■ **Disabled**: Transactions are complete as soon as the writes are cached in memory. When this policy is active, completion of transactions does not depend on waiting for writes to reach the disk. This setting affects performance, scalability, and reliability. **Note:** The use of the direct-write policy is reliable in Solaris systems, but Windows systems may leave transaction data in the on-disk cache without writing it to disk immediately. This is considered unreliable, because a power failure can cause loss of on-disk cache data, possibly resulting in lost or duplicate messages. For reliable writes using the direct-write policy on Windows, either disable all write caching for the disk (enabled by default), or use a disk with a battery-backed cache. Some file systems, however, do not allow this value to be changed (for example, a RAID system that has a reliable cache). **Note:** If the JMS file store is used exclusively for paging non-persistent messages to the disk, the synchronous write policy is ignored.

3. After updating the settings, click **Next**.

   Review the domain settings and proceed with domain creation as described in Section 4.12, "Review the Domain Settings and Create the Domain."

For more information, see *Configuring WebLogic Server Environments*.

6 Extending WebLogic Domains

You can add product component functionality, or additional applications and services to an existing domain by extending it using the Configuration Wizard. For example, if you created a base WebLogic Server domain and you want to add the Avitek Medical Records Examples functionality, you must extend the domain by using the `medrec.jar` template.

The Configuration Wizard simplifies the task of extending an existing domain by using extension templates. For more information about domain templates, see *WebLogic Server Domain Template Reference*.

Figure 6–1 shows the steps in the process of extending a domain.

Figure 6–1 Domain Extension Process

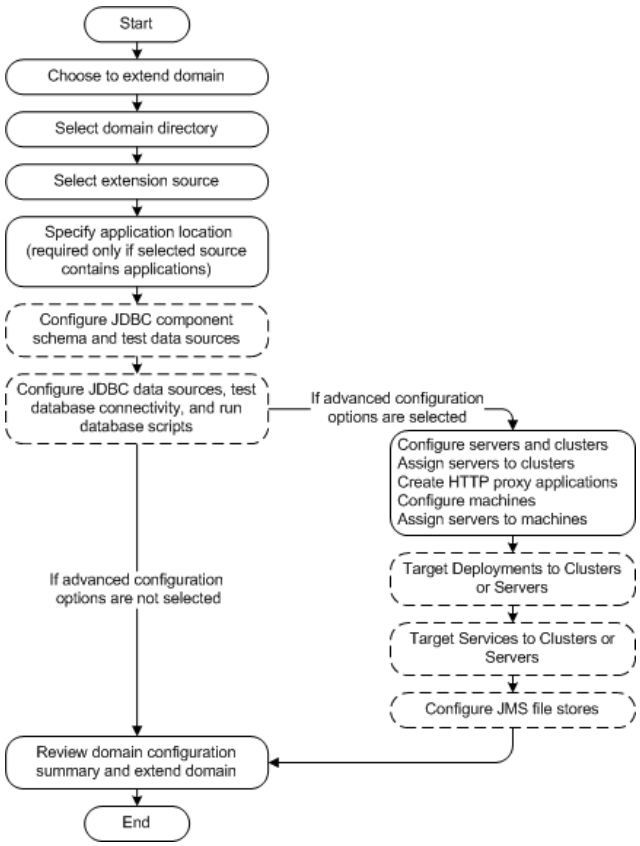

Note: When you extend a domain in which the managed servers are distributed (by using the `pack` and `unpack` commands) to remote machines, the startup scripts on the remote machines are not updated automatically.

To ensure that startup scripts on the remote machines are updated, perform the following steps after extending the domain.

1. Delete the managed server directories on the remote machines.
2. Create a managed server template from the extended domain, by using the `pack` command with the `-managed=true` option.
3. Create managed servers on the remote machines, by `unpack`ing the managed server template.

For more information about the `pack` and `unpack` commands, see *Oracle Fusion Middleware Creating Templates and Domains Using the Pack and Unpack Commands*.

To extend a domain, perform the following steps:

1. Ensure that the servers in the domain that you want to extend are not running.
2. Start the Configuration Wizard as described in Section 2.1, "Starting the Configuration Wizard in Graphical Mode."

   The **Welcome** screen of the Configuration Wizard prompts you to choose whether you want to create a domain or extend an existing domain by adding product component functionality, applications and services.
3. Select **Extend an existing WebLogic domain** and click **Next**.

   The **Select a WebLogic Domain Directory** screen prompts you to select the existing domain to update with additional applications or services.
4. Use the navigation tree to select a valid domain directory (a directory that contains a `config.xml` file in the `config` directory of the domain, and click **Next**.

   The **Select Extension Source** screen prompts you to select the source from which to extend your domain. You can select products to add to your domain, or extend your domain by using an existing extension template.
5. Choose one of the following options:

   - **Extend my domain automatically to support the following added products:**

     Select the check boxes for the products to add to your domain. The products already included in your domain are indicated by grayed-out check boxes.

   - **Extend my domain using an existing extension template**

     Specify the path to the extension template in the **Template location** field.

   Click **Next**.

   If you opted to extend the domain by using an extension template, and if the template that you specified contains applications, the **Specify Domain Name and Location** screen is displayed.
6. Specify the path to the application in the **Applications location** field, and click **Next**.

The remaining steps in the process to extend domains are identical to the steps for creating a domain, except that the wizard does not display the screens to configure the administration server and RDBMS security store database.

7. Perform the advanced configuration tasks (such as configuring servers, clusters, and machines; and targeting deployments and services to clusters or servers), as described in Chapter 4, "Customizing the Domain Environment."
8. Configure the JDBC component schema and JDBC data sources as described in Chapter 5, "Customizing JDBC Data Sources and Component Schema."
9. Finally, in the **Configuration Summary**, screen review the configuration changes and then click **Extend**.

   The **Creating Domain** screen displays status messages during the domain extension process. The Configuration Wizard updates the `config.xml` file and other application-specific components in the domain directory, as defined by the domain template.
10. Click **Done**.