



BEA WebLogic Enterprise Security™®

Administration Application Installation

Product Version: 4.2 Service Pack 2
Document Revised: September 18, 2005

Copyright

Copyright © 2005 BEA Systems, Inc. All Rights Reserved.

Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms of the BEA Systems License Agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form without prior consent, in writing, from BEA Systems, Inc.

Use, duplication or disclosure by the U.S. Government is subject to restrictions set forth in the BEA Systems License Agreement and in subparagraph (c)(1) of the Commercial Computer Software--Restricted Rights Clause at FAR 52.227-19; subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, subparagraph (d) of the Commercial Computer Software--Licensing clause at NASA FAR supplement 16-52.227-86; or their equivalent.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA Systems DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

Trademarks or Service Marks

BEA, BEA JRocket, BEA Liquid Data for WebLogic, BEA WebLogic Server, Built on BEA, Jolt, JoltBeans, SteelThread, Top End, Tuxedo, and WebLogic are registered trademarks of BEA Systems, Inc. BEA AquaLogic, BEA AquaLogic Data Services Platform, BEA AquaLogic Enterprise Security, BEA AquaLogic Service Bus, BEA AquaLogic Service Registry, BEA Builder, BEA Campaign Manager for WebLogic, BEA eLink, BEA Manager, BEA MessageQ, BEA WebLogic Commerce Server, BEA WebLogic Enterprise, BEA WebLogic Enterprise Platform, BEA WebLogic Enterprise Security, BEA WebLogic Express, BEA WebLogic Integration, BEA WebLogic Java Adapter for Mainframe, BEA WebLogic JDriver, BEA WebLogic JRocket, BEA WebLogic Log Central, BEA WebLogic Personalization Server, BEA WebLogic Platform, BEA WebLogic Portal, BEA WebLogic Server Process Edition, BEA WebLogic WorkGroup Edition, BEA WebLogic Workshop, and Liquid Computing are trademarks of BEA Systems, Inc. BEA Mission Critical Support is a service mark of BEA Systems, Inc. All other company and product names may be the subject of intellectual property rights reserved by third parties.

All other trademarks are the property of their respective companies.

Contents

1. Overview

About this Document	1-1
Prerequisites for This Guide	1-2
Product Documentation on the dev2dev Web Site	1-2
Related Information	1-3
Product Overview	1-4
Audience	1-4
Security Environment	1-4
Security Architecture Functional Description	1-6

2. Preparing to Install

Installation and Distribution	2-1
Web Distribution	2-2
CD-ROM Distribution	2-2
Installation Prerequisites	2-3
System Requirements	2-3
Licensing	2-5
Requirements for Reinstalling the Administration Application	2-6
Selecting Directories for the Installation	2-6
BEA Home Directory	2-6
Product Installation Directory	2-7

3. Setting Up and Administering the Database

Setting Up and Administering the Oracle Database and Client	3-2
Before you Begin the Oracle Database Setup	3-2
Overview of the Oracle Client/Server Architecture	3-2
Oracle Database System Requirements	3-4
Installing and Configuring the Oracle Database	3-5
Installing the Oracle Database	3-5
Creating an Instance of an Oracle Database	3-8
Configuring an Oracle Policy Database	3-9
Installing and Configuring an Oracle Client	3-10
Installing and Configuring an Oracle Client on Windows	3-11
Installing and Configuring the Oracle Client on Sun Solaris	3-14
Installing and Configuring the Oracle Client on Red Hat Advanced Server 2.13-15	
Installing and Configuring the Oracle Client on Red Hat Advanced Server 3.03-18	
Tuning an Oracle Database	3-21
Calculating Oracle Tablespace Requirements	3-22
Calculating Oracle Tablespace Size Requirements	3-23
Calculating Oracle Rollback Tablespace Size Requirements	3-24
Optimizing the Oracle Database for Large Policies	3-26
Administering an Oracle Policy Database	3-26
Creating a User Account in an Oracle Policy Database	3-28
Using the Database Administration Utilities with Oracle	3-30
Backing Up an Oracle Database	3-31
Setting Up and Administering the Sybase Database and Client	3-32
Before you Begin the Sybase Database Setup	3-32
Overview of the Sybase Client/Server Architecture	3-32
Sybase Database System Requirements	3-34

Installing and Configuring the Sybase Adaptive Server	3-35
Installing the Sybase Database	3-35
Setting the Sybase Database Configuration Parameters	3-38
Creating Sybase Database Devices	3-39
Creating and Configuring a Sybase Policy Database	3-39
Installing and Configuring a Sybase Database Client	3-42
Testing an Existing Sybase Open Client Installation	3-42
Installing and Configuring the Sybase Open Client on Windows	3-43
Installing and Configuring the Sybase Open Client on Sun Solaris	3-44
Installing and Configuring the Sybase Open Client on Red Hat Advanced Server 2.1	3-45
Tuning the Sybase Database	3-46
Calculating Sybase Database Size Requirements	3-46
Calculating Sybase Tablespace Requirements	3-47
Calculating Sybase Data Size Requirements	3-48
Calculating Sybase Transaction Log Size Requirements	3-49
Preventing Database Log Bloat with Sybase	3-50
Expanding the Policy Database with Sybase	3-50
Optimizing the Sybase Database for Large Policies	3-50
Administering the Sybase Policy Database	3-51
Creating a User Account in a Sybase Policy Database	3-52
Using the Database Administration Utilities with Sybase	3-53
Backing Up a Sybase Database	3-55

4. Installing

Before you Begin	4-1
System Security and BEA Weblogic Enterprise Security	4-2
System Users	4-2

System Groups	4-2
File System Permissions	4-3
Secure User Names and Passwords	4-3
Generating a Verbose Installation Log	4-6
Starting the Installation Program on Windows Platforms	4-7
Starting the Installation Program on a Sun Solaris Platform	4-8
Starting the Installation Program on a Linux Platform	4-11
Running the Installation Program	4-12
What's Next	4-19
Installing a Secondary Administration Application	4-19

5. Post Installation Tasks

Installing the Policy Database Schema	5-2
Installing the Policy Database Schema on Windows	5-3
Installing the Policy Database Schema on Sun Solaris	5-4
Installing the Policy Database Schema on Linux	5-5
Starting and Stopping Processes	5-6
Logging into the Administration Console	5-6
Changing the System Password	5-7
Fine Tuning your Application	5-8
What's Next?	5-9

6. Configuring Metadirectories

Why Use Metadirectories?	6-2
Metadirectory Configuration Overview	6-3
Preparing to Configure a Metadirectory	6-4
Installing the RadiantOne Synchronization Services	6-5
Configuring Metadirectory Tables and Database Triggers	6-5

Creating Metadirectory Destination Tables	6-5
Metadirectory Destination Table Guidelines and Restrictions	6-6
Creating Metadirectory Destination Tables Using Oracle or Sybase	6-8
Configuring a JDBC Connection Pool and JMS	6-9
Configuring Metadirectory Database Triggers.	6-11
Configuring Metadirectory Schemas.	6-13
Extracting the Source Schemas	6-14
Loading the Source Schemas	6-17
Extracting the Destination Schemas.	6-19
Loading the Destination Schemas	6-22
Configuring the Source-to-Destination Topology	6-23
Configuring the Topology Transformations.	6-25
UpLoading User and Group Data.	6-26
Configuring Metadirectory Synchronization.	6-28
Configuring the Synchronization Hub	6-30
Configuring the Directory Connector.	6-30
Configuring the Policy Database Connectors.	6-31
Starting the Synchronization Hub	6-33
Starting the Source and Destination Connectors	6-33
Verifying that Metadirectory Synchronization Works	6-34

7. Uninstalling

Uninstalling the Administration Application on Windows	7-1
Additional Steps	7-2
Uninstalling the Administration Application on Solaris or Linux.	7-3

Index

Overview

BEA WebLogic Enterprise Security is certified by standard BEA platform requirements. The Administration Application is an *n*-tier Java-based web application that is primarily based on JavaServer Pages (JSP). The Administration Application includes a variety of utilities, including the Administration Console, Service Control Manager, Policy Importer, and Policy Exporter. The console allows you to manage and configure services and policies for any number of distributed Security Service Modules. This guide provides the information needed to install the Administration Application, including system requirements and prerequisite software and hardware. It does not include information for additional Security Service Modules that you may also be installing.

This section covers the following topics:

- [“About this Document” on page 1-1](#)
- [“Product Overview” on page 1-4](#)

About this Document

This document provides application developers with the information needed to setup the database, install the BEA WebLogic Enterprise Security™ Administration Application, and configure metadirectories. The document is organized as follows:

- [Chapter 1, “Overview,”](#) provides an introduction to the Administration Application and the WebLogic Enterprise Security component architecture.
- [Chapter 2, “Preparing to Install,”](#) discusses system requirements (software and hardware) that you need to ensure are met before installing the Administration Application.

- [Chapter 3, “Setting Up and Administering the Database,”](#) explains how to configure your Oracle or Sybase client and database server, which is the repository for all policy data.
- [Chapter 4, “Installing,”](#) provides detailed procedures for installing the BEA WebLogic Enterprise Security Administration Application.
- [Chapter 5, “Post Installation Tasks,”](#) provides detailed procedures for tasks you need to perform after the installation, including loading the policy schema into the database, and starting and stopping services.
- [Chapter 6, “Configuring Metadirectories,”](#) describes how to configure your metadirectories and how to create the synchronization service used by the Administration Console to synchronize user data stored in the policy database to user data stored in an external repository. This step is optional.
- [Chapter 7, “Uninstalling,”](#) describes the procedures for uninstalling the BEA WebLogic Enterprise Security Administration Application.

Prerequisites for This Guide

Prior to reading this guide, you should read the [Introduction to BEA WebLogic Enterprise Security](#). This document describes how the product works and provides conceptual information that is helpful to understanding the necessary installation components.

Additionally, BEA WebLogic Enterprise Security includes many unique terms and concepts that you need to understand. These terms and concepts—which you will encounter throughout the documentation—are defined in the [Glossary](#).

Product Documentation on the dev2dev Web Site

BEA product documentation, along with other information about BEA software, is available from the BEA dev2dev web site:

<http://dev2dev.bea.com>

To view the documentation for a particular product, select that product from the Product Centers menu on the left side of the screen on the dev2dev page. Select More Product Centers. From the BEA Products list, choose WebLogic Enterprise Security 4.2. The home page for this product is displayed. From the Resources menu, choose Documentation 4.2. The home page for the complete documentation set for the product and release you have selected is displayed.

Related Information

The BEA corporate web site provides all documentation for BEA WebLogic Enterprise Security. Other BEA WebLogic Enterprise Security documents that may be of interest to the reader include:

- *Introduction to BEA WebLogic Enterprise Security*—This document provides overview, conceptual, and architectural information for the Administration Application products.
- *BEA WebLogic Enterprise Security Administration Guide*—This document provides a complete overview of the product and includes step-by-step instructions on how to perform various administrative tasks.
- *BEA WebLogic Enterprise Security Policy Managers Guide*—This document defines the policy model used by BEA WebLogic Enterprise Security, and describes how to generate, import and export policy data.
- *Programming Security for Java Applications*—This document describes how to implement security in Java applications. It includes descriptions of the security service Application Programming Interfaces and programming instructions.
- *Programming Security for Web Services*—This document describes how to implement security in web servers. It includes descriptions of the Web Services Application Programming Interfaces.
- *Developing Security Providers for BEA WebLogic Enterprise Security*—This document provides security vendors and security and application developers with the information needed to develop custom security providers.
- *Javadocs for Java API*—This document provides reference documentation for the Java Application Programming Interfaces that are provided with and supported by this release of BEA WebLogic Enterprise Security.
- *Wsdl docs for Web Services API*—This document provides reference documentation for the Web Services Application Programming Interfaces that are provided with and supported by this release of BEA WebLogic Enterprise Security.
- *Javadocs for Security Service Provider Interfaces*—This document provides reference documentation for the Security Service Provider Interfaces that are provided with and supported by this release of BEA WebLogic Enterprise Security.
- *Javadocs for BLM API*—This document provides reference documentation for the Business Logic Manager (BLM) Application Programming Interfaces that are provided with and supported by this release of BEA WebLogic Enterprise Security.

Product Overview

The following sections describe the audience for this document and provide an overview of the Administration Application:

- “Audience” on page 1-4
- “Security Environment” on page 1-4
- “Security Architecture Functional Description” on page 1-6

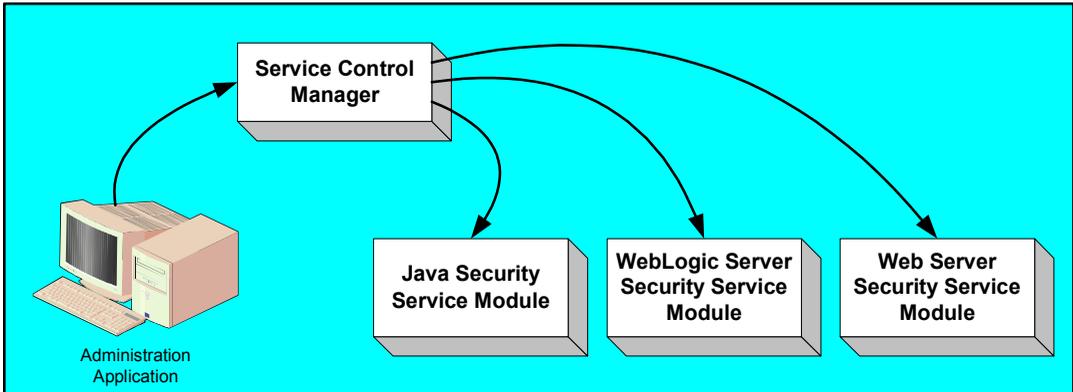
Audience

It is assumed that readers understand web technologies and have a general understanding of the Microsoft Windows or UNIX operating system being used. The general audience for this installation guide includes Database Administrators and System Administrators.

Security Environment

Figure 1-1 shows the BEA WebLogic Enterprise Security system environment.

Figure 1-1 WebLogic Enterprise Security Environment



- Administration Application

The Administration Application allows you to manage and configure multiple Security Service Modules. While Security Service Modules specify and consume configuration data and then service security requests accordingly, the Administration Application allows you to configure and display the security providers that are deployed in the Security Service Module, and to modify the configuration data for those providers.

- Service Control Manager

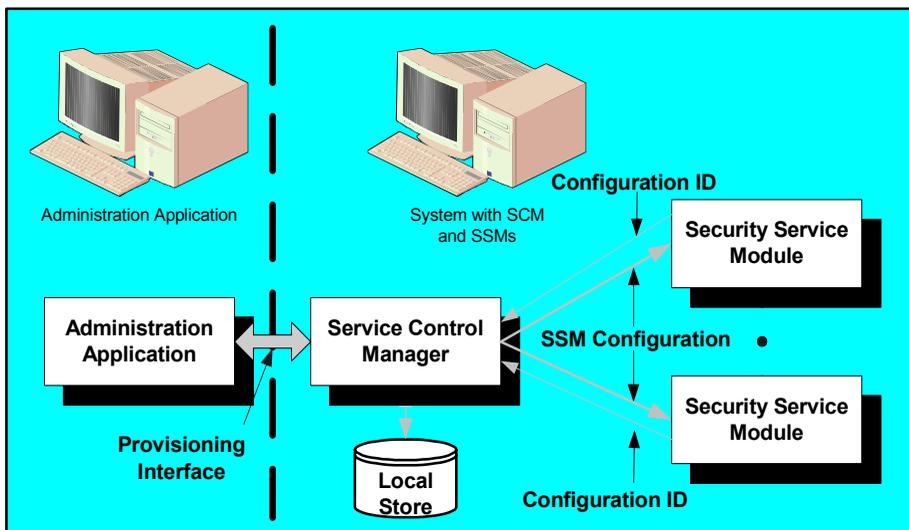
The Service Control Manager is an essential component of the BEA WebLogic Enterprise Security provisioning mechanism and is a key component of a fully-distributed security enforcement architecture. A Service Control Manager is a machine agent that exposes a provisioning interface to the Administration Application to facilitate the management of a potentially large number of distributed Security Service Modules. A Service Control Manager can receive and store meta-data updates, both full and incremental, initiated by the Administration Application.

The Administration Application uses the provisioning mechanism of the Service Control Manager to distribute configuration and policy data to each Security Service Module where it is stored locally (see [Figure 1-2](#)). Security Service Modules can be distributed throughout an enterprise and can be embedded in Java applications, application servers, and web servers.

- Security Service Module

After you use the Administration Application to configure an instance of a Security Service Module with configuration and policy data, the Security Service Module does not require any additional communication with the Service Control Manager to provide security services. However, the Service Control Manager maintains communication with the Security Service Module to distribute full and incremental updates whenever necessary.

Figure 1-2 Deploying Configuration and Policy Data

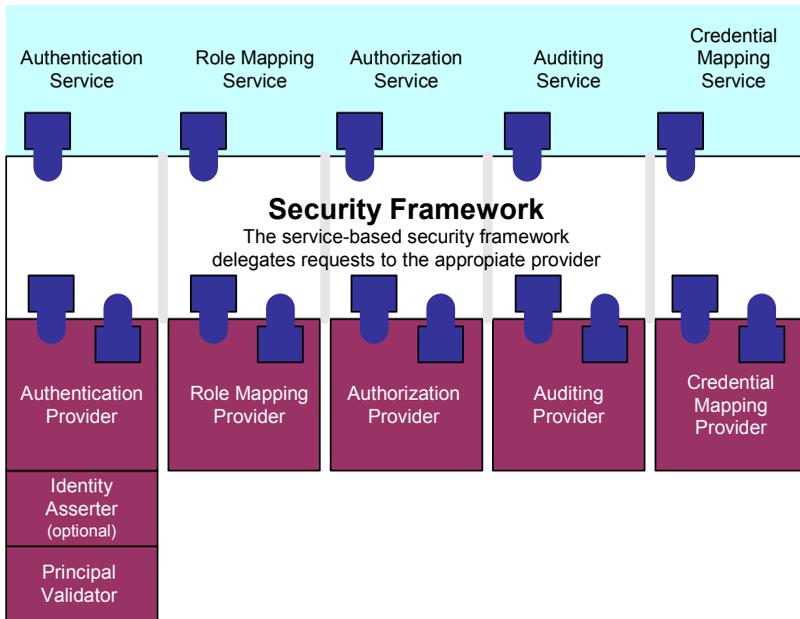


Security Architecture Functional Description

Figure 1-3 shows the major architectural components of the Administration Application. The following topics describe these components:

- “Security Services” on page 1-6
- “Security Framework” on page 1-7
- “Security Providers” on page 1-7

Figure 1-3 Architectural Components



Security Services

The Administration Application supports the following security services:

- Authentication Service
- Role Mapping Service
- Authorization Service
- Auditing Service

- Credential Mapping Service

Security Framework

The primary function of the Security Framework is to provide an application programming interface (API) that security and application developers can use to implement security functions. Within that context, the Security Framework also acts as an intermediary between security services and the security providers that are configured into the Security Service Module.

Security Providers

When you install the Administration Application or a Security Service Module, a JAR file is deployed that contains all the security providers that ship with the product. However, before any of the security providers can be used, you must configure them through the Administration Application. You have the option of configuring either the security providers that ship with the product or custom security providers, that you develop or purchase from third-party security vendors. You can configure your providers through the Administration Console or by importing policy datafiles directly into the database. The Administration Application supports the following types of security providers:

- Authentication provider
- Identity Assertion provider
- Credential Mapping provider
- Role Mapping provider
- Authorization provider
- Auditing provider

Note: To use custom security providers with a Security Service Module, you must deploy the security provider MBean JAR file (MJF) to both the provider directory on the machine on which you install the Security Service Module and on the Administration Server. For more information on how to develop custom security providers, see [Developing Security Providers for BEA WebLogic Enterprise Security](#).

For more information on security providers, see [Introduction to BEA WebLogic Enterprise Security](#).

Overview

Preparing to Install

This section provides the information needed to install the BEA WebLogic Enterprise Security Administration Application, including system requirements, and prerequisite software and hardware. It does not include information for installing a Security Service Module.

This section covers the following topics:

- [“Installation and Distribution” on page 2-1](#)
- [“Installation Prerequisites” on page 2-3](#)
- [“Selecting Directories for the Installation” on page 2-6](#)

Installation and Distribution

BEA WebLogic Enterprise Security products are distributed and installed using the BEA Installation and Distribution System, which provides a complete framework for the following:

- Distribution of BEA products by download from the BEA web site
- Installation and uninstallation of the BEA WebLogic Enterprise Security Administration Application including documentation

BEA WebLogic Enterprise Security is distributed on both the BEA web site and on CD-ROM.

Web Distribution

If you want to install the product by downloading it from the BEA web site, contact BEA Sales at <http://www.bea.com/framework.jsp?CNT=sales1.htm&FP=/content/about/contact/> and request a download.

The package installer downloads a stand-alone version of the installation program that contains the complete Administration Application. The package installer is approximately 70 MB.

Documentation is available from the product documentation home page. Be sure to download the most up-to-date information from the BEA web site at:

<http://e-docs.bea.com/wles/docs42/download.html>.

CD-ROM Distribution

If you purchased BEA WebLogic Enterprise Security from your local sales representative, you will find the following items in the product box:

Four CD-ROMs:

- Disk 1 of 4 contains the following BEA WebLogic Enterprise Security products:
 - Administration Application software for Microsoft Windows platforms
 - Security Service Modules software for Microsoft Windows platforms
 - Documentation in both PDF and HTML format
- Disk 2 of 4 contains the following BEA WebLogic Enterprise Security products:
 - Administration Application software for Linux and Sun Solaris
 - Security Service Modules software for Linux and Sun Solaris
- Disk 3 of 4 contains the BEA WebLogic Enterprise Security metadirectory software for Microsoft Windows platforms. This product is used with the Administration Application to integrate user repositories.
- Disk 4 of 4 contains the BEA WebLogic Enterprise Security metadirectory software for Linux and Sun Solaris platforms.

The following printed documents:

- Introduction to BEA WebLogic Enterprise Security
- BEA Software License and Limited Warranty pamphlet

- Customer Support Quick Reference and Other Important Information card

Installation Prerequisites

The Administration Application requires certain software components to operate properly. Review these requirements carefully before installing the product.

- “System Requirements” on page 2-3
- “Licensing” on page 2-5
- “Requirements for Reinstalling the Administration Application” on page 2-6

System Requirements

Table 2-1 lists the system requirements for the Administration Application.

Note: The machine on which you install the Administration Application must have a static IP address. The IP address is used by the Security Service Modules and Service Control Manager for connectivity. Also, on a Windows platform, the file system must be configured for NTFS and not FAT.

Table 2-1 System Requirements

Use	Component and Version
WebLogic Server 8.1 with Service Service Pack 3 or 4	WebLogic Enterprise Security Administration Server requires that you install the BEA WebLogic Server in the BEA_HOME directory. You can download WebLogic Server from: http://commerce.bea.com/showallversions.jsp?family=WLS
Java Development Kit (JDK)	The installation program requires JRE 1.4.2_04 or JRE 1.4.2_05, which are installed as part WebLogic Server 8.1, SP3 and SP4 respectively. Note: The installation process sets JAVA_HOME and related variables to point to this directory. All scripts installed use JAVA_HOME by default.

Table 2-1 System Requirements (Continued)

Use	Component and Version
Database Storage	<ul style="list-style-type: none"> • Oracle 8i (8.1.7) or Oracle 9i Release 2 (9.2.0.1) or • Sybase Adaptive Server Enterprise, Version 12.5 <p>Note: BEA recommends using the Oracle 9.2.0.4 client on Linux platforms. Use of an earlier version may seriously increase the amount of system memory used by the WebLogic Enterprise Security servers or processes. This behavior can eventually cause the server to use up the system memory. The use of Oracle 9.2.0.4 does not exhibit this behavior.</p> <p>For instructions on how to set up your database, see Chapter 3, “Setting Up and Administering the Database.”</p>
Database Connectivity	<p>Oracle or Sybase Client Runtime Software</p> <p>BEA recommends that the version of your client software be the same as the database to which you are connecting. Do not use an older version of the client software to connect to a newer version of the database server.</p> <p>For instructions on how to set up the database client, see “Installing and Configuring an Oracle Client” on page 3-10 and “Installing and Configuring a Sybase Database Client” on page 3-42.</p>
Platforms supported	<p>The BEA WebLogic Enterprise Security Administration Application runs on any of the following platforms:</p> <ul style="list-style-type: none"> • Intel Pentium compatible with Microsoft Windows NT 4.0 • Intel Pentium compatible with Microsoft Windows 2000 Professional • Intel Pentium compatible with Microsoft Windows 2000 Server/Advanced Server • Sun Microsystems Sparc with Solaris, versions 8 and 9 • Linux Red Hat Advanced Server, version 3.0
Web Browser	<p>Microsoft Internet Explorer, Version 5.5 or later; 6.0 is recommended. In addition, the Java Plug-in for Internet Explorer from the Java Runtime Environment (JRE) 1.4 or greater is required.</p>
Display Resolution	<p>A display resolution of 1024 x 768 or higher is recommended when running the Administration Console.</p>

Table 2-1 System Requirements (Continued)

Use	Component and Version
Memory	256 MB of RAM minimum, 512 MB or more is recommended. Each user session requires approximately 5 MB of memory.
Hard Disk Space	<p>About 100 MB free storage space for the installed product (this does not include WebLogic Server storage space).</p> <p>Refer to the database installation instructions for recommendations on database storage allocation.</p>
Certificates and Keystores	<p>BEA WebLogic Enterprise Security uses an implementation of the Transport Layer Security (TLS) 1.0 specification (see TLS Protocol). The server hosting the WebLogic Enterprise Security Administration Application supports TLS on a dedicated listen port that defaults to 7010. To establish a secure connection, a Web browser connects to the Administration Server by supplying the listen port and the secure address (HTTPS) in the connection URL, for example, https://myserver:7010.</p>
User Repository	<p>Optionally, an external directory server or database, containing your user store, configured through the BEA WebLogic Enterprise Security Metadirectory.</p> <p>This product is included with your software. You must install it according to the instructions, and then configure it to provision the Administration Application metadirectory. For further information on how to use metadirectory with the Administration Application, see “Configuring Metadirectories” on page 6-1.</p>
Reporting	<p>Optionally, you can use Log4j to configure a reporting application to support auditing features. For further information on how to use Log4j with the Administration Application, see: http://jakarta.apache.org/log4j/docs/.</p>

Licensing

The product software cannot be used without a valid license. When you install the Administration Application, the installation program creates an evaluation license that expires in 90 days.

To use the Administration Application in a production environment, you must purchase a license. For information about purchasing a license, contact your BEA Sales Representative.

Requirements for Reinstalling the Administration Application

If you are installing the Administration Application on a computer on which the Administration Application was previously installed, refer to [“Uninstalling” on page 7-1](#) and make sure all of the uninstall steps were completed; otherwise the installation may fail.

Selecting Directories for the Installation

During installation, you need to specify locations for the following directories:

- [“BEA Home Directory” on page 2-6](#)
- [“Product Installation Directory” on page 2-7](#)

BEA Home Directory

During installation, you are prompted to choose an existing BEA Home (`BEA_HOME`) directory. You should specify the same BEA Home directory that you specified when you installed WebLogic Server 8.1. The BEA Home directory is a repository for common files that are used by multiple BEA products installed on the same machine. For this reason, the BEA Home directory can be considered a "central support directory" for the BEA products installed on your system. The files in the BEA Home directory are essential to ensuring that BEA software operates correctly on your system. They perform the following types of functions:

- Ensure that licensing works correctly for the installed BEA products
- Facilitate checking of cross-product dependencies during installation

The files and directories in the BEA Home (`BEA_HOME`) directory are described in your WebLogic documentation. Although it is possible to create more than one BEA Home directory, BEA recommends that you avoid doing so. In almost all situations, a single BEA Home directory is sufficient. There may be circumstances, however, in which you prefer to maintain separate development and production environments on a single machine, each containing a separate product stack. With two directories, you can update your development environment (in a BEA Home directory) without modifying the production environment until you are ready to do so.

Product Installation Directory

The product installation directory contains all the software components used to administer BEA WebLogic Enterprise Security. During installation, you are prompted to choose a product installation directory. If you accept the default, the software is installed in the following directory:

`c:\bea\wles42-admin` (Windows)

`/opt/bea/wles42-admin` (Sun Solaris and Linux)

where `c:\bea` is the `BEA_HOME` directory and `wles42-admin` is the product installation directory. You can specify any name and location on your system for your product installation directory and there is no requirement that you name the directory `wles42-admin` or create it under the BEA Home directory.

Preparing to Install

Setting Up and Administering the Database

This section provides information and guidelines to assist you in installing, configuring, and managing the database server and the database client to used with the WebLogic Enterprise Security Administration Application. This information is not meant to replace or supersede in any way the database documentation provided by Oracle and Sybase for their database server and client products. Also, the information provided here assumes that you are familiar with the Oracle database documentation.

BEA WebLogic Enterprise Security stores all policy and configuration data used by the Administration Application and Security Service Modules in the policy database. You can use either an Oracle database or a Sybase database for your policy data storage. You must install and configure the database server software before you install the Administration Application. If you install the Administration Application on a machine other than the machine on which you install the database, you must also install and configure the respective Oracle or Sybase client on that machine.

Note: To perform a database installation and setup, you must be a database administrator with a database administrator username and password and permission to create a new instance. In addition, you should be knowledgeable about the operating system you are working with and be adept at database installations and configuration issues. If you do not feel comfortable performing any of these tasks, ask your database administrator for assistance.

This section covers the following topics:

- [“Setting Up and Administering the Oracle Database and Client” on page 3-2](#)
- [“Setting Up and Administering the Sybase Database and Client” on page 3-32](#)

Setting Up and Administering the Oracle Database and Client

This section contains the procedures for setting up and administering an Oracle database and an Oracle Client. It covers the following topics:

- [“Before you Begin the Oracle Database Setup” on page 3-2](#)
- [“Installing and Configuring the Oracle Database” on page 3-5](#)
- [“Installing and Configuring an Oracle Client” on page 3-10](#)
- [“Tuning an Oracle Database” on page 3-21](#)
- [“Administering an Oracle Policy Database” on page 3-26](#)

Before you Begin the Oracle Database Setup

Before you install and set up your Oracle database, review the following topics to better understand Oracle database configuration requirements:

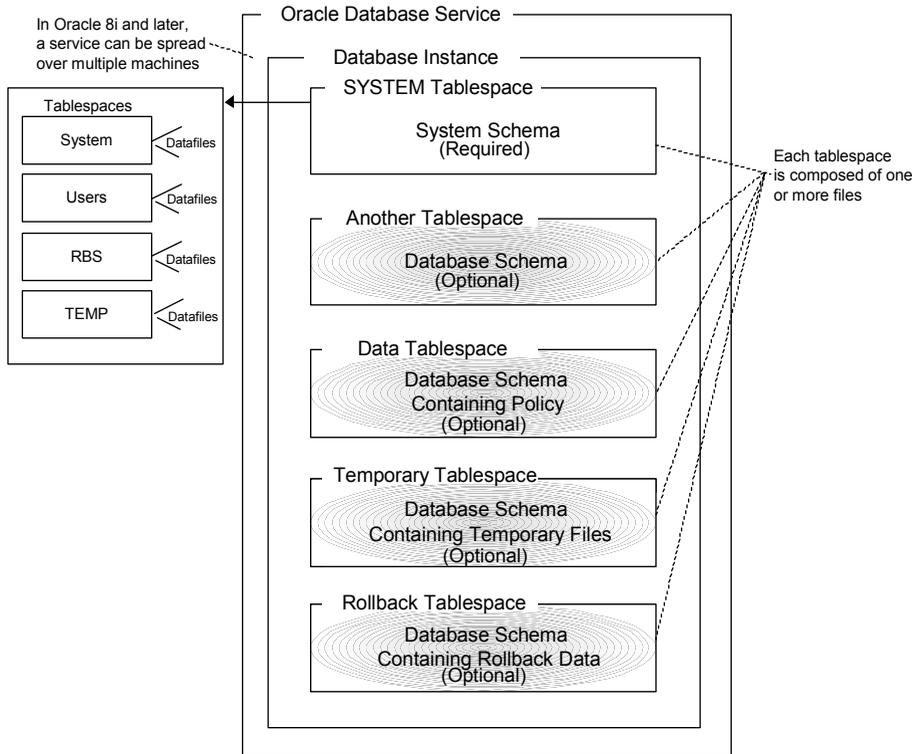
- [“Overview of the Oracle Client/Server Architecture” on page 3-2](#)
- [“Oracle Database System Requirements” on page 3-4](#)

Overview of the Oracle Client/Server Architecture

The Oracle database service is the server in the Oracle client/server architecture (see [Figure 3-1](#)). The database service manages a database instance and multiple database users, keeps track of the actual location of data on disks, maintains mapping of logical data to physical data storage, and maintains data and procedure caches in memory. In this section, the example of the Oracle service name is viewed from the client perspective.

Each Oracle service is identified by a global database name and an Oracle system identifier referred to as the `SID`. The Oracle global database name is the full name of a database that uniquely differentiates it from any other databases in your network domain. One global database name can represent several database instances. The global database name is also known as the service name. The `SID` distinguishes the database instance from any other database instances on the same machine.

Figure 3-1 Oracle Database Setup



An Oracle instance is a running Oracle database made up of memory structures and background processes. Each instance is associated with an `SID`. With the Oracle Parallel Server, multiple instances can exist on different machines for a single database.

The policy database is a set of database schemas in which all data are stored. A database schema is a collection of objects associated with a particular schema name. The objects include tables, views, domains, constraints, assertions, privileges, and so on.

A datafile is an Oracle term for a file that contains the contents of logical database structures, such as tables and indexes. One or more datafiles form a logical unit of storage called a tablespace. A datafile is associated with only one tablespace and only one database.

A tablespace is a logical portion of a database used to allocate storage for table and index data. Each tablespace corresponds to one or more physical datafiles. Every Oracle database has a tablespace called `SYSTEM` and may have additional tablespaces. A tablespace is used to group

related logical structures. The database username or user ID is a login that is given permission by the database administrator to access a specific database instance. This user is also called the schema owner, that is, the owner of the schema objects such as tables, views and triggers that are created.

Oracle Database System Requirements

Table 3-1 describes the minimum requirements for the system on which the Oracle database server is installed.

Table 3-1 Oracle Setup Requirements

Requirement	Description
Software version	<p>Oracle database server:</p> <ul style="list-style-type: none"> Version 8i (8.1.7) Version 9i (9.0.1) and 9i Release 2 (9.2.0.1) <p>Note: On Linux platforms, BEA recommends using the Oracle 9.2.0.4 client. Use of an earlier version may seriously increase the amount of system memory used by the WebLogic Enterprise Security servers or processes. This behavior can eventually cause the server to use up the system memory. The use of 9.2.0.4 does not exhibit this behavior.</p>
Server platform	Any platform supported by Oracle.
Memory	As required by Oracle server installation (64 MB minimum).
Disk space for the starter database	As required by Oracle server installation, plus space required to store policy data; 500 MB recommended.
Disk space for Oracle software	Refer to your installation guide for the Oracle Database Server.
Disk space for policy database	Minimum of one tablespace with 250 MB of free space is required. To approximate space requirements for any policy size, use the formula in “Calculating Oracle Tablespace Size Requirements” on page 3-23.
Oracle Client	Oracle Client that ships with your version of the product. BEA requires that the version of your client software be the same as the database to which you are connecting. Do not use an older version of the client software to connect to a newer version of the database server.

Installing and Configuring the Oracle Database

This section provides additional instructions for installing and configuring an Oracle database for use with the WebLogic Enterprise Security Administration Application.

To install and configure the database, perform the following tasks:

- [“Installing the Oracle Database” on page 3-5](#)—Use this procedure only if you are going to install the Oracle database software and create and configure an instance of the database.
- [“Creating an Instance of an Oracle Database” on page 3-8](#)—Use this procedure only if the Oracle database software is already installed and you want to create another instance of the database.
- [“Configuring an Oracle Policy Database” on page 3-9](#)—Use this procedure to configure a policy database for use by the WebLogic Enterprise Security Administration Application.

Installing the Oracle Database

This section provides recommendations for installing the Oracle database and creating a database instance. When you run the Oracle installation program, it automatically starts the Database Configuration Assistant, which you use to create an instance of the database. If the Oracle database is already installed on the database host machine, you can skip this procedure and go to [“Creating an Instance of an Oracle Database” on page 3-8](#) and then go to [“Configuring an Oracle Policy Database” on page 3-9](#).

To install the Oracle database and create a database instance, perform these steps:

1. Install the Oracle database according to instructions in the *Oracle Database Installation Guide* and system requirements defined in [Table 3-1](#). During the installation, define the following parameters.
 - Global Database Name—The full Oracle database name that distinguishes the database from any other databases in your network domain.
 - Database System Identifier—The Oracle system identifier (SID). The SID distinguishes the database instance from any other database instances on the same machine.
 - SYS and SYSTEM Passwords—The Oracle database install program creates two user accounts, SYS and SYSTEM, and assigns default passwords. During the installation, you are prompted to change these passwords. For security reasons, Oracle recommends that you specify new passwords for these user accounts when you install the database software.

Note: Be sure to record the settings you use for these parameters, because you will need them later in this procedure and also to configure the Oracle Client if you are required to do so.

2. Use the Oracle Database Configuration Assistant to configure the database. When the Database Configuration Assistant starts, step through the screens and select the Template Name and Memory settings as specified in [Table 3-2](#).

Note: The Memory setting only applies to Oracle 9i databases.

Table 3-2 Database Configuration Assistant Settings

Database Configuration Assistant Setting	Recommended Setting
Template Name	General Purpose. Note: This selection specifies the template to use to create the instance of the database.
Memory (Typical or Custom), for 9i only	Custom Note: This selection is on the Initialization Parameters screen.

3. Use one of the following procedures to set the database initialization parameters.
 - a. For Oracle 8i, open the `ora.init` file located in `ORACLE_HOME/admin/db_name/pfile` directory and go to the step 4.
 - b. For Oracle 9i, click All Initialization Parameters on the Initialization Parameters screen of the Database Configuration Assistant. The All Initialization Parameters screen appears (see [Figure 3-2](#)). Go to step 4.

Figure 3-2 Oracle Initialization Parameters Screen

Name	Value	Included (Y/N)	Category
optimizer_features_enable	9.0.1		Optimizer
remote_dependencies_m...	TIMESTAMP		PL/SQL
parallel_threads_per_cpu	2		Parallel Executions
logmnr_max_persitent_s...	1		Miscellaneous
nls_date_language			NLS
workarea_size_policy	MANUAL		Sort, Hash Joins, Bitmap
O7_DICTIONARY_ACCES...	FALSE		Security and Auditing
license_max_sessions	0		License Limits
star_transformation_enabl...	FALSE	✓	Optimizer
nls_date_format			NLS
lock_sga	FALSE		SGA Memory
fixed_date			Miscellaneous

4. Refer to [Table 3-3](#) and enter the initialization parameters.

Table 3-3 Initialization Parameters Recommended Values

Database Parameter Name	Recommended Value
shared_pool_size	6925926 Bytes This value must be large enough for good server performance.
db_block_buffers	1000
log_buffer	32768 Bytes
processes	150
db_block_size	8192 Bytes (or greater) Note: Block Size is critical. Some Oracle installs set this value to 4096 by default, which creates problems for some scripts. You must set this value to 8192 or larger.

Table 3-3 Initialization Parameters Recommended Values (Continued)

Database Parameter Name	Recommended Value
open_cursors	500
rollback_segments	80. See “Calculating Oracle Rollback Tablespace Size Requirements” on page 3-24.

- Proceed through the Oracle Database Configuration Assistant pages to the Creation Options page, select Create Database, and click Finish. The assistant creates an instance of the Oracle database.
- Set your system PATH environment variables as shown in [Listing 3-1](#).

Note: In [Listing 3-1](#), <drive> is the hard drive on which the Oracle database is installed and <version> is either 90 or 92.

Listing 3-1 System PATH Environment Variable Settings on Windows

```

For Oracle 9i:
<drive>:\oracle\ora<version>\bin;
C:\Program Files\Oracle\jre\1.3.1\bin;
C:\Program Files\Oracle\jre\1.1.8\bin;

For Oracle 8i:
<drive>:\oracle\ora81\bin;
C:\Program Files\Oracle\jre\1.1.7\bin;
    
```

- Proceed to “Configuring an Oracle Policy Database” on page 3-9.

Creating an Instance of an Oracle Database

This section describes how to create and configure an instance of an Oracle database. It assumes that the Oracle database software was installed.

Note: You should only perform this procedure when you want to create and configure instances of the database in addition to the instance that was created when the database software was installed.

Perform the following steps to create an instance of an Oracle database:

Note: The section provides guidance to assist you, but it does not supersede the documentation provided by Oracle.

1. To start the Oracle Database Configuration Assistant, click Start>Programs>Oracle-<*OraHome*>Configuration and Migration Tools>Database Configuration Assistant, where *OraHome* indicates the version of the software. The Database Configuration Assistant starts.
2. When the Database Configuration Assistant starts, step through the screens and select Template Name and Memory settings as specified in [Table 3-2](#).

Note: The Memory setting only applies to Oracle 9i databases.
3. Use one of the following procedures to set the initialization parameters.
 - a. For Oracle 8i, open the `ora.init` file that is located in `ORACLE_HOME/admin/db_name/pfile` directory and go to the step 4.
 - b. For Oracle 9i, click the All Initialization Parameters bottom on the Initialization Parameters screen of the Database Configuration Assistant. The All Initialization Parameters screen appears (see [Figure 3-2](#)). Go to step 4.
4. Refer to [Table 3-3](#) and enter the database configuration parameters.
5. Proceed through the Oracle Database Configuration Assistant pages to the Creation Options page, select Create Database, and click Finish. The assistant creates an instance of the Oracle database.
6. To configure a policy database for this instance of an Oracle database, see [“Configuring an Oracle Policy Database” on page 3-9](#).

Configuring an Oracle Policy Database

To configure an Oracle policy database, you must create the policy database, create a security role and a user, and grant the security role and user access.

To configure a policy database, perform the following steps:

1. Open a command window, run the Oracle SQLPlus utility, and login as user `SYSTEM` with the password you set for that user account when you installed the Oracle database software.

```
sqlplus system/password@asi
```

where: *password* is the password you set for the system account when you installed the database software and *asi* is the database instance name.

2. To configure the policy database, enter the following commands at the SQL> prompt:

```
SQL>connect sys as sysdba
SQL>create tablespace DATA datafile 'C:/Oracle/oradata/ASI/data.dbf'
      size 10M autoextend on next 1M MAXSIZE 250M;
SQL>CREATE ROLE asi_role;
SQL>GRANT CREATE SESSION to asi_role;
SQL>GRANT CREATE TABLE to asi_role;
SQL>GRANT CREATE PROCEDURE to asi_role;
SQL>GRANT CREATE SEQUENCE to asi_role;
SQL>GRANT CREATE TRIGGER to asi_role;
SQL>GRANT CREATE VIEW to asi_role;
SQL>CREATE USER wles IDENTIFIED BY password
      default tablespace DATA QUOTA UNLIMITED on DATA;
SQL>GRANT asi_role to wles;
SQL>GRANT SELECT on SYS.V_$LOCKED_OBJECT to wles;
```

where: *asi_role* is the security role you define, *wles* is the user you define, and *password* is the user password.

3. To verify that the configured user can connect to the policy database, open a command window and type:

```
sqlplus wles/password@asi
```

where: *wles* and *password* are the user and *password* you defined and *asi* is the database instance name.

This completes the configuration of the instance of the policy database.

Installing and Configuring an Oracle Client

If you intend to install the WebLogic Enterprise Security Administration Application on the same machine as you installed the Oracle database, you do not need to install or configure the Oracle Client. The Oracle database installation includes the Oracle Client, so you can skip this task.

However, if you intend to install the Administration Application on a machine other than the machine on which the Oracle database is installed, you must install and configure an Oracle client on that machine to be able to access the Oracle database server from the client machine.

To install and configure an Oracle Client, you need to know the following information:

- The Global Database Name that you defined when you created the database instance.
- The host name of the database machine
- The port number on which the database instance is running (the default port number is 1521).

- The policy database username and password that you defined when you configured the policy database.

For instructions on installing and configuring an Oracle Client, see the following topics:

- [“Installing and Configuring an Oracle Client on Windows”](#) on page 3-11
- [“Installing and Configuring the Oracle Client on Sun Solaris”](#) on page 3-14
- [“Installing and Configuring the Oracle Client on Red Hat Advanced Server 2.1”](#) on page 3-15
- [“Installing and Configuring the Oracle Client on Red Hat Advanced Server 3.0”](#) on page 3-18

Installing and Configuring an Oracle Client on Windows

To install and configure an Oracle Client, perform these steps:

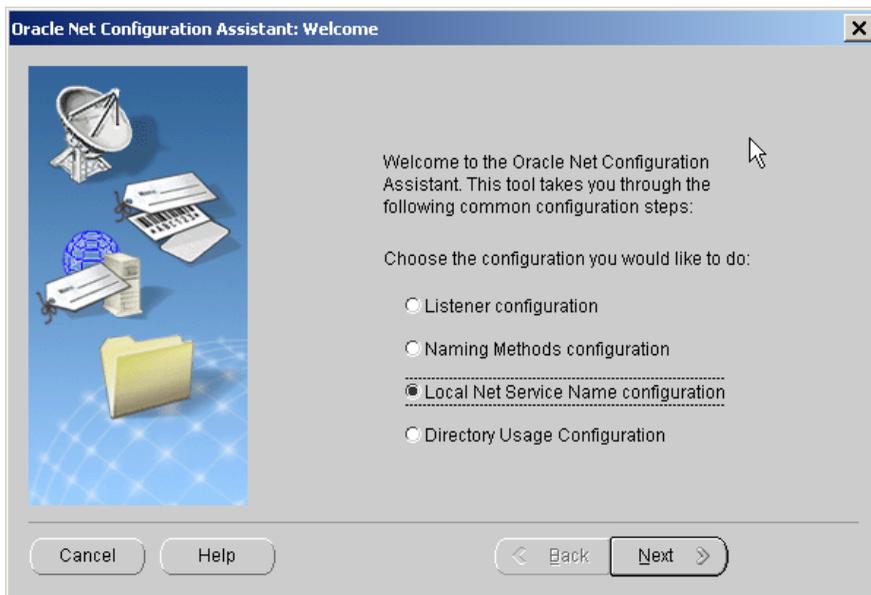
Note: This section provides guidance to assist you, but it does not supersede the documentation provided by Oracle.

1. Install the Oracle Client according to instructions in the *Oracle Database Installation Guide* for Windows. If the Oracle Client is already installed, skip this step and go to the next step.
2. Start the Oracle Net Configuration Assistant and use it to configure a Local Net Service Name entry for connecting to the Oracle database instance (see [Figure 3-3](#)).

Note: [Figure 3-3](#) shows the Oracle 9i screen. The Oracle 8i screen offers the same options.

In this step, you set up a service entry in the Oracle configuration file, which is located on the client machine at: `ORACLE_HOME/network/admin/tnsnames.ora`.

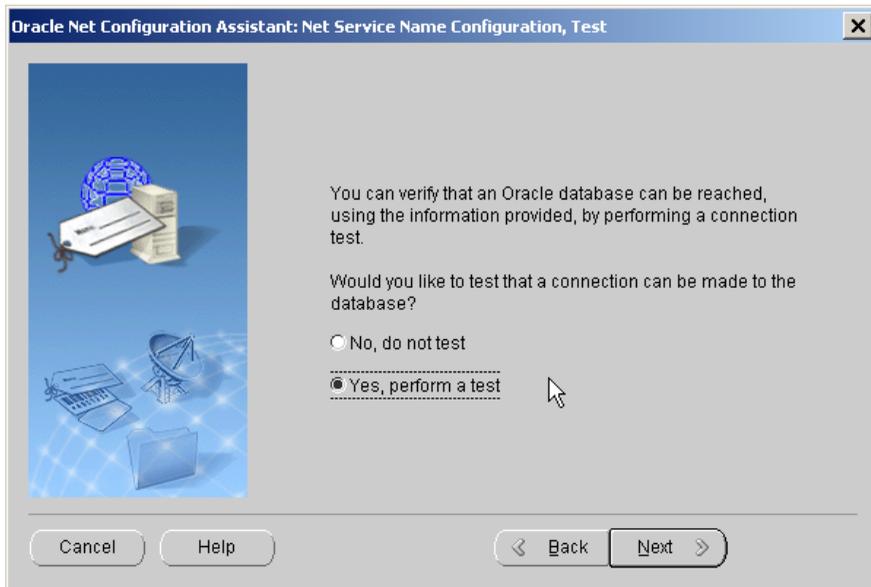
Figure 3-3 Oracle Net Configuration Assistant: Welcome Page (Oracle 9i)



3. To verify that the Oracle Client can access the Oracle database, at the Net Configuration Assistant screen (see [Figure 3-4](#)), select the **Yes, perform a test** radio button, click **Next**, and execute the test.

Note: [Figure 3-4](#) shows the Oracle 9i screen. The Oracle 8i screen offers the same options.

Figure 3-4 Oracle Net Service Name Configuration Test Page (Oracle 9i)



4. If the test in the previous step fails, click the Change Login button on the test results page, enter the database username and password, and execute the test again.

Note: If you experience problems getting the Oracle Client to connect to the Oracle database instance, check the configuration of the database instance in the `ORACLE_HOME\ora<version>\network\admin\tnsnames.ora` file located on the database server host machine, where `<version>` is 81, 90, or 92.

5. To use SQLplus to connect to the Oracle database instance on the machine on which your Oracle client is running as the `wles` user, open a command window and type:

```
sqlplus wles/password@asi
```

where: `wles` and `password` are the user and password you defined when you configured the policy database and `asi` is the database instance name.

This completes the configuration of the Oracle Client.

Installing and Configuring the Oracle Client on Sun Solaris

To install and configure the Oracle Client on a Sun Solaris platform, perform these steps:

Note: This section provides guidance to assist you, but it does not supersede the documentation provided by Oracle.

1. If they do not already exist, have a Sun Solaris system administrator create a group called `dba` and a user ID called `oracle`.
2. Set `dba` as the primary group for `oracle`.
3. Log into Sun Solaris as `oracle`.
4. Unload the Oracle client software to a local directory using the Oracle Installer.
5. Set the `ORACLE_HOME` environment variable to the local directory. If necessary, refer to your Oracle Installation Guide.
6. Set the `PATH` environment variable to include the `bin` subdirectory of `$ORACLE_HOME`.
7. Set the `LD_LIBRARY_PATH` environment variable to include the `lib` subdirectory of `$ORACLE_HOME`.
8. To connect to the Oracle database instance on the machine on which your Oracle client is running, open a command window and type the following SQLplus command:

```
sqlplus wles/password@asi
```

where: `wles` and `password` are the user and password you defined when you configured the policy database and `asi` is the database instance name.

If this command is successful, the client is configured, and you can skip the next step of this procedure. If this command fails, proceed to step 9.

9. Start an Oracle Network Configuration tool, such as Net Configuration Assistant or Net Manager, and configure a local net service name entry for connecting to the database instance. This step sets up a service entry in the Oracle configuration file located at: `$ORACLE_HOME/network/admin/tnsnames.ora`.

Note: You may also use a text editor to edit the `tnsnames.ora` file. However, you should be familiar with Oracle Net before editing the `tnsnames.ora` file with a text editor.

This completes the configuration of an Oracle Client.

Installing and Configuring the Oracle Client on Red Hat Advanced Server 2.1

There may be some additional considerations when installing Oracle 9i and 8i Clients on Red Hat Advanced Server 2.1. To understand all the considerations relative to installing on the Red Hat Advanced Server in your environment, see the Oracle and Red Hat documentation.

Note: If you are installing the Oracle 8i Client on Red Hat Advanced Server 2.1, the Net8 Configuration tool may hang during the installation process. To start the Net8 Configuration tool, you need to download and install JRE-1.1.8v3, and switch the JRE to use the proper version of the tool: `$ORACLE_HOME/bin/netasst`, by changing the value for `JREDIR`.

To install and configure an Oracle Client on Red Hat Advanced Server 2.1, perform the following steps:

Note: This section provides guidance to assist you, but it does not supersede the documentation provided by Oracle.

1. If you are installing by downloading the software from the Oracle web site, go to step 2. If you are installing from an Oracle CD-ROM, skip step 2, and go to step 3.
2. Using the instructions provided on the Oracle download site, perform the following steps:
 - a. Download the Oracle Database Server software from the Oracle web site. For example, the Oracle 9.2 download kit requires that you download the following files:

```
ship_9204_linux_disk1.cpio.gz
ship_9204_linux_disk2.cpio.gz
ship_9204_linux_disk3.cpio.gz
```

- b. To unzip each file, run:

```
gunzip <filename>
```

- c. To extract the `cpio` archive, run the following command on each file:

```
cpio -idmv <filename>.cpio
```

This command creates directories named `Disk1`, `Disk2`, and `Disk3`.

Note: The Oracle 8i Database Server software is no longer available from the Oracle download site. The Oracle 8i Database Server is now in Extended Support (ES) mode.

3. To start the Oracle installer, run the following command from `Disk1`:

```
./runInstaller
```

4. Select the Oracle Client for installation, and then select the Administrative edition or Application Programmer edition.
5. When an error window appears, wait for the following error message:

```
Error in invoking target install of makefile
/path/app/oracle/product/version/xyz/lib/ins_xyz.mk, and prompt for Retry,
Ignore, and Cancel, where xyz may be precomp, or plssql, or something else and
version is either 8.1.7 or 9i.
```

6. When this error occurs, examine the file: \$ORACLE_HOME/install/make.log.

The file contains the following lines of text.

```
path/app/oracle/product/version/bin/genclntsh
/lib/libc.so.6: undefined reference to \Q_dl_lazy@GLIBC_2.1.1'
/lib/libc.so.6: undefined reference to \Q_dl_dst_substitute@GLIBC_2.1.1'
/lib/libc.so.6: undefined reference to \Q_dl_out_of_memory@GLIBC_2.2'
/lib/libc.so.6: undefined reference to \Q_dl_relocate_object@GLIBC_2.0'
/lib/libc.so.6: undefined reference to \Q_dl_clktck@GLIBC_2.2'
/lib/libc.so.6: undefined reference to \Q__libc_enable_secure@GLIBC_2.0'
/lib/libc.so.6: undefined reference to \Q_dl_catch_error@GLIBC_2.0'
.....
/usr/bin/ld: cannot find -lclntsh
collect2: ld returned 1 exit status
/bin/chmod: getting attributes of \Qprocob18': No such file or directory
make: *** [procob18] Error 1
/usr/bin/make -f ins_precomp.mk relink
ORACLE_HOME=/pathora/u01/app/oracle/product/version EXENAME=ott...
```

7. Set the environment variables for ORACLE_HOME, PATH and LD_LIBRARY_PATH.
8. Open another window, and change to the \$ORACLE_HOME/bin directory.
9. Edit the genclntsh script by setting LD_SELF_CONTAINED="".
10. Run the following command:

```
./genclntsh
```

The following message appears:

```
Created /path/app/oracle/product/version/lib/libclntst#.a
```

11. Return to the Oracle installer, and click Retry.

12. After linking the Oracle libraries, the installer prompts you to run `root.sh`.

Note: Before continuing with step 13, for Oracle 8i, edit `root.sh` and change the line:
`RMF=/bin/rm -f` to `RMF="/bin/rm -f"`, and the line that starts with `RUID=...` by
adding a single quote just before the last back-slash (`\`).

13. Log in as root and run:

```
./root.sh
```

14. Return to the installer, and click OK to continue.

The installer continues. At the last step, it starts the Net Configuration tool to let you configure the first Net Service Name.

15. To connect to the Oracle database instance on the machine on which your Oracle client is running, open a command window and type the following SQLplus command:

```
sqlplus wles/password@asi
```

where: `wles` and `password` are the user and password you defined when you configured the policy database and `asi` is the database instance name.

If this command is successful, the client is configured and you can skip the remaining steps of this procedure. If this command fails, proceed to step 16.

16. Use the Net Configuration Assistant to configure a local net service name entry for connecting to the database instance. This step sets up a service entry in the Oracle configuration file (`$ORACLE_HOME/network/admin/tnsnames.ora`).

Note: If you installing the Oracle 8i Client on Red Hat Advanced Server 2.1, the Net8 Configuration tool may hang during the installation process. Abort that process. To start the Net8 Configuration tool, download JRE-1.1.8v3, and switch the JRE to use the proper version of the tool: `$ORACLE_HOME/bin/netasst`, by changing the value for `JREDIR`.

17. Exit the installer.

Note: If you installing the Oracle 8i Client on Red Hat Advanced Server 2.1, apply the client patch: `glibc-2.1.3-stubs.tar.gz` that you downloaded earlier.

This completes the configuration of an Oracle Client.

Installing and Configuring the Oracle Client on Red Hat Advanced Server 3.0

There may be some additional considerations when installing Oracle 9i and 8i Clients on Red Hat Advanced Server 3. To understand all the considerations relative to installing on the Red Hat Advanced Server in your environment, see the Oracle and Red Hat documentation.

Note: If you are installing the Oracle 8i Client on Red Hat Advanced Server 3.0, the Net8 Configuration tool may hang during the installation process. To start the Net8 Configuration tool, you need to download and install JRE-1.1.8v3, and switch the JRE to use the proper version of the tool: `$ORACLE_HOME/bin/netasst`, by changing the value for `JREDIR`.

To install and configure an Oracle 9.2 Client on Red Hat Advanced Server 3.0, perform the following steps:

Note: This section provides guidance to assist you, but it does not supersede the documentation provided by Oracle.

1. If you are installing by downloading the software from the Oracle web site, go to step 2. If you are installing from an Oracle CD-ROM, skip step 2, and go to step 3.
2. Using the instructions provided on the Oracle download site, perform the following steps:
 - a. Download the Oracle Database Server software from the Oracle web site. For example, the Oracle 9.2 download kit requires that you download the following files:

```
ship_9204_linux_disk1.cpio.gz  
ship_9204_linux_disk2.cpio.gz  
ship_9204_linux_disk3.cpio.gz
```

- b. To unzip each file, run:

```
gunzip <filename>
```

- c. To extract the `cpio` archive, run the following command on each file:

```
cpio -idmv <filename>.cpio
```

This command creates directories named `Disk1`, `Disk2`, and `Disk3`.

Note: The Oracle 8i Database Server software is no longer available from the Oracle download site. The Oracle 8i Database Server is now in Extended Support (ES) mode.

3. Set the environment variable `LD_ASSUME_KERNEL` to `2.4.1`.
4. Install the following RedHat Package Managers (RPMs):

```
compat-db-4.0.14-5.i386.rpm \  
compat-gcc-7.3-2.96.122.i386.rpm \  
compat-gcc-7.3-2.96.122.i386.rpm \
```

```
compat-gcc-c++-7.3-2.96.122.i386.rpm \
compat-libstdc++-7.3-2.96.122.i386.rpm \
compat-libstdc++-devel-7.3-2.96.122.i386.rpm \
```

5. Relink `gcc` to `gcc296` and `g++` to `g++296`.

Note: Be sure to restore the `gcc` and `g++` to `gcc323` and `g++323` after the installation.

6. Download the patch `p3006854_9204_LINUX.zip` from <http://metalink.oracle.com/>. For more information, see Oracle bug 3006854. To apply this patch, run:

```
su - root
# unzip p3006854_9204_LINUX.zip
Archive:  p3006854_9204_LINUX.zip
  creating: 3006854/
  inflating: 3006854/rhel3_pre_install.sh
  inflating: 3006854/README.txt
# cd 3006854
# sh rhel3_pre_install.sh
Applying patch...
Patch successfully applied
```

7. Go to the `Disk1` directory and run this command: `./runInstaller`.

Note: You cannot run this command as root.

Note: If you are accessing the system through a Telnet connection, make sure that your display is set correctly.

The `./runInstaller` command displays the Oracle Universal Installer: Welcome window.

8. On the Oracle Universal Installer Welcome window, click Next. The Inventory Location window appears.
9. On the Inventory Location window, set the directory field to where you want to install Oracle, for example: `/export/home/oracle`. The UNIX Group Name window appears.
10. On the UNIX Group Name window, enter the name for your group, and click Next.
11. A message window opens and directs you to run the `/tmp/orainstRoot.sh` command as root. Running this command outputs the following two lines:

```
Creating Oracle Inventory pointer file (/etc/oraInst.loc)
Changing groupname of /export/home/oracle to engineering.
```

12. Return to the message window and click Continue. The File Locations window appears.

13. On File Locations window, verify that the Source field is correct and change the Destination Name and Path to where you want to store the oracle files, and click Next. For example:

```
Name: ORACLE
Path: /export/home/oracle
```

The Loading products progress indicator displays in the upper right corner of the window. When the loading completes, the Available Products window appears.

14. On the Available Products window, select Oracle 9i Client 9.2.0.1.0, and click Next. The Installation Types window appears.
15. On the Installation Types window, select the Runtime radio button and click Next. The Summary window appears.
16. On the Summary window, click Install. The Install window appears and a progress indicator displays showing the status of the installation process. When the installation completes, the following message is displayed:

```
A configuration script needs to be run as root before installation can proceed. Please leave this window up, run /export/home/oracle/root.sh as root from another window, then come back here and click OK to continue.
```

17. Run the root.sh command. The root.sh command outputs the following:

```
Running Oracle9 root.sh script...
\nThe following environment variables are set as:
ORACLE_OWNER= dboot
    ORACLE_HOME= /export/home/oracle
Enter the full pathname of the local bin directory: [/usr/local/bin]:
Copying dbhome to /usr/local/bin ...
Copying oraenv to /usr/local/bin ...
Copying coraenv to /usr/local/bin ...
\nCreating /etc/oratab file...
Adding entry to /etc/oratab file...
Entries will be added to the /etc/oratab file as needed by Database
Configuration Assistant when a database is created
Finished running generic part of root.sh script.
Now product-specific root actions will be performed.
```

18. After the script completes, click OK. The Configuration Tools window appears. Click No on the Oracle Net Configuration Assistant: Welcome window, and click Next.
19. Select the oracle8i or later database or service radio button on the Oracle Net Configuration Assistant: Net Service Name Configuration, Database Version window, and click Next.

20. Enter a Service Name into the entry field, and click Next. For example:
`mydbhost.mydomain.com`.
21. Select TCP on the Oracle Net Configuration Assistant: Net Service Name Configuration. Select the Protocols window, and click Next.
22. Enter a host name into the entry field on the Oracle Net Configuration Assistant: Net Service Name Configuration, TCP/IP Protocol window, and click Next. For example:
`mydbhost.mydomain.com`.
23. Select Yes to perform a test on the Oracle Net Configuration Assistant: Net Service Name Configuration Test window, and click Next. You should get this message:
`Connecting...Test successful.`
If not, click Back, correct the settings, and retest. If successful, click Next.
24. Enter a Net Service Name value on the Oracle Net Configuration Assistant: Net Service Name Configuration Net Service Name window, and click Next. For example: `mydbhost`.
25. Select No on the ...Another Net Service Name window, and click Next.
26. Click Next on the ...Configuration Done window, and click Next.
27. Click Finish to complete the Configuration process.
28. On the Oracle Universal Installer: End of Installation window, click Exit to close the Oracle installation.

This completes the configuration of an Oracle Client.

Tuning an Oracle Database

After you have installed and configured the Oracle database and the Oracle Client, you should tune the database to suit the needs of your particular environment. The following topics provide information to assist in tuning your Oracle database:

- [“Calculating Oracle Tablespace Requirements” on page 3-22](#)
- [“Calculating Oracle Tablespace Size Requirements” on page 3-23](#)
- [“Calculating Oracle Rollback Tablespace Size Requirements” on page 3-24](#)
- [“Optimizing the Oracle Database for Large Policies” on page 3-26](#)

Calculating Oracle Tablespace Requirements

To determine the tablespace size requirements, allot the amount of disk space based on the size of your policy. You should use 250 MB as an absolute minimum, provided the rollback segments can handle the policy loading and distribution.

To determine your actual tablespace requirements, see the following topics:

- [“Minimum Disk Space Allotment” on page 3-22](#)
- [“Group Flattening and Rules” on page 3-22](#)
- [“Metadirectory Synchronization Services” on page 3-23](#)

Minimum Disk Space Allotment

The 250 MB minimum disk-space allotment works fine with a small policy and a small user community such as the following:

- The policy has a maximum of 1000 users
- Each user belongs to no more than one group
- Each user has one single-valued attribute
- The policy has less than 100 privileges, resources, and declarations
- The policy has less than 100 flattened rules; no composite privileges, resources, or subjects (users and groups) in the rule

Group Flattening and Rules

Group flattening means that a rule can exist in one of two forms: a simple rule or a composite rule. A composite rule is a combination of two or more simple rules to make them easier to use. The process for reducing a composite rule to its component simple rules is called "flattening the group."

For example, if you had three local users named Joe, Betty, and Sam, you could grant those users a role in an application by creating a composite rule like this:

```
Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,  
      [//user/AcctDept/Joe/, //user/AcctDept/Sam/, //user/AcctDept/Betty/]);
```

In the policy language, this rule means "grant Joe, Sam, and Betty, who belong to the `AcctDept`, the role of `bookkeeper` in the accounting application, `AcctApp`."

The rule is a composite rule because it reduces or flattens to these three simple rules:

```
Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,
//user/AcctDept/Joe/);

Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,
//user/AcctDept/Sam/);

Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,
//user/AcctDept/Betty/);
```

Even though you may see one composite rule, the composite is actually stored and distributed as three flattened simple rules. The main ramification of rule flattening is that your policies can take much more disk space than you might think when simply looking at your policy. For information on how to construct rules, see [Securing Resources and Defining Policy Rules](#) in the *Policy Managers Guide*.

Metadirectory Synchronization Services

If you want to use the BEA WebLogic Enterprise Security Metadirectory Synchronization Services, you must create an additional set of tables to use to synchronize identity information. As a result, the amount of space required to store identity information approximately doubles so allocate an appropriate amount of extra tablespace. For more information, see [Configuring Metadirectories](#) in *WebLogic Enterprise Security Administration Application Installation*.

Calculating Oracle Tablespace Size Requirements

You can estimate your space requirements using the following formulas. With group flattening, as with rules, group memberships are also reduced or flattened to their simple data components. For example, if you have a user that belongs to a group through group inheritance, the membership is stored as though the user were a direct member of the group. Thus, there is a separate group to user mapping for each group in the inheritance hierarchy. All numeric results are represented in megabytes. All formulas use the variables described in [Table 3-4](#).

Table 3-4 Oracle Variables

Variable	Description
a	Total number of user attribute values for all users, in thousands
d	Total number of declarations, in thousands
m	Total number of flattened user/group mappings, in thousands
o	Total number of objects, in thousands
p	Total number of privileges, in thousands

Table 3-4 Oracle Variables (Continued)

Variable	Description
q	Total number of object attribute values for all resources, in thousands
r	Total number of flattened rules, in thousands
u	Total number of users, in thousands

Oracle Corporation recommends using multiple datafiles for any tablespace that approaches one GB in size.

Use the following formula to calculate your tablespace size requirements. For a description of the formula variables, see [Table 3-4](#).

$$\text{Data Tablespace} = 250 + 0.3u + 0.2 a + 0.1m + 1.2(o + p) + 0.75(q-1) + 4d + 5r$$

For example, if all the variables had the value 5, the formula looks like this:

$$= 250 + 0.3(5) + 0.2(5) + 0.1(5) + 1.2(5 + 5) + 0.75(5-1) + 4(5) + 5(5)$$

and reduces to this:

$$= 250 + 1.5 + 1 + 0.5 + 12 + 3 + 20 + 25$$

and finally:

$$= 313$$

Thus, the example requires a minimum of 313 MB of disk space.

Calculating Oracle Rollback Tablespace Size Requirements

The rollback tablespace is required to successfully distribute the largest policy changes between distributions. When you change the policy and distribute it frequently in smaller chunks, the space required is reduced dramatically.

$$\text{Rollback Tablespace} = 250 + 2.5(o + p) + 2.5(q-1) + 6d + 10r$$

For a very small policy (the built-in policy plus a few hundred users), you can use the system rollback segments that are created during the database installation. However, BEA recommends that you create a new tablespace with a few rollback segments. Configuring 250 MB of rollback segments works fine for the restricted policy mentioned earlier.

For more information on configuring tablespace requirements, see the following topics:

- [“Temporary Tablespace Requirements” on page 3-25](#)
- [“Adding Additional Tablespaces” on page 3-25](#)
- [“Creating the Rollback Segments” on page 3-26](#)

Temporary Tablespace Requirements

For a very small policy (the built-in policy plus a few hundred users), you can use the system temporary tablespace (`TEMP`) that is created during the database installation. For larger policy, check to ensure that your `TEMP` setting is sufficient. However, BEA recommends that you create a new temporary tablespace that is at least one-fourth the size of your data tablespace.

Adding Additional Tablespaces

The datafile name and tablespace sizes in the following instructions are given for illustration purposes only. You should determine your own needs and replace these values. In addition, BEA chose to use the `autoextend` option in the instructions, but your needs may differ. Consult your Oracle documentation for details.

Finally, the following instructions are specific to a Sun Solaris installation. If you are installing on Windows 2000, replace all the forward slashes with back slashes and begin all file paths with the drive name.

To add additional tablespaces, perform the following steps:

1. To login as the system administrator, open a command window and type:

```
sqlplus SYSTEM/password@asi
```

where: *password* is the password you defined when you installed the database software
asi is the database instance name.

2. To create the data tablespace, at the sqlplus prompt, type:

```
SQL> create tablespace DATA datafile '/oradata/ASI/data.dbf' size 10 M
      autoextend on next 1M MAXSIZE 250M;
```

where: `DATA` is the tablespace name and `/oradata/ASI/data.dbf` is the physical datafile used to store the database schema.

3. To create the rollback tablespace, type:

```
SQL> Create tablespace RBS datafile '/oradata/ASI/rbs.dbf' size 10 M
      autoextend on next 1M MAXSIZE 250M;
```

where: `RBS` is the tablespace name and `/oradata/ASI/rbs.dbf` is physical datafile to contain the rollback schema.

Creating the Rollback Segments

Use the instructions provided in this section to create and enable the maximum number of rollback segments (five) in the rollback tablespace created previously. You may want to do this if the rollback segments for the default database installation are not sufficient. Depending on the size of the rollback tablespace (represented in the commands as `rbs_1` to `rbs_5`), you can either create and enable more segments or increase the size of the existing segments instead.

To create the rollback segments, open a command window, start SQLplus, and type the following commands:

```
SQL> create rollback segment rbs_1 tablespace RBS STORAGE(INITIAL 100K
      NEXT 100K OPTIMAL 500K MINEXTENTS 2 MAXEXTENTS 100);
SQL> create rollback segment rbs_2 tablespace RBS STORAGE(INITIAL 100K
      NEXT 100K OPTIMAL 500K MINEXTENTS 2 MAXEXTENTS 100);
SQL> create rollback segment rbs_3 tablespace RBS STORAGE(INITIAL 100K
      NEXT 100K OPTIMAL 500K MINEXTENTS 2 MAXEXTENTS 100);
SQL> create rollback segment rbs_4 tablespace RBS STORAGE(INITIAL 100K
      NEXT 100K OPTIMAL 500K MINEXTENTS 2 MAXEXTENTS 100);
SQL> create rollback segment rbs_5 tablespace RBS STORAGE(INITIAL 100K
      NEXT 100K OPTIMAL 500K MINEXTENTS 2 MAXEXTENTS 100);
```

Optimizing the Oracle Database for Large Policies

When your Oracle database contains a large policy, you may want to do one or more of the following to optimize performance:

- Ensure that you allot the maximum amount of RAM to the Oracle server in the initialization parameters file (for Oracle 9i, `SPFILESID.ORA`; for Oracle 8i, `INIT.ORA`).
- Ensure that you increase `SORT_AREA_SIZE` for the Oracle server in the initialization parameters file.
- Ensure that you allot enough disk space for the data tablespace and rollback segments.
- Run `install_sort_oracle.bat` or `install_sort_oracle.sh` to install ASCII sorting, instead of the default dictionary sorting that comes with the database schema installation. This improves the Administration Console response time. See [“Administering an Oracle Policy Database” on page 3-26](#) for details.

Administering an Oracle Policy Database

This section covers the following topics:

- [“Creating a User Account in an Oracle Policy Database” on page 3-28](#)
- [“Using the Database Administration Utilities with Oracle” on page 3-30](#)
- [“Backing Up an Oracle Database” on page 3-31](#)

Creating a User Account in an Oracle Policy Database

This section describes how to configure a new user account in an Oracle policy database. This account is necessary so that the policy for the instance of the Administration Application managed by this user can have a dedicated storage area allocated in the database instance.

Note: To perform this procedure, you must log into the Oracle database server as a database administrator.

To set up a database user account, perform these steps:

1. To login to the Oracle database server, type:

```
sqlplus dba/password@ASERVER
```

where:

dba is the username you use to access the database.

password is your database administrator password.

ASERVER is the name of the Oracle service (as defined in your *tnsnames.ora* file).

2. To create a new role in the database server, type:

```
SQL> create role asi_role;  
SQL> grant create session to asi_role;  
SQL> grant create table to asi_role;  
SQL> grant create procedure to asi_role;  
SQL> grant create sequence to asi_role;  
SQL> grant create trigger to asi_role;  
SQL> grant create view to asi_role;
```

where: *asi_role* is the new role.

The following example uses the default tablespaces generated when the Oracle database was first installed, although you can specify any tablespaces.

3. To set up a new database user account, type:

Note: In this example, you use the default tablespaces generated when you created and configured the Oracle database instance, however, you can specify any tablespaces.

```
SQL> create user username identified by password  
SQL> default tablespace users quota unlimited on users  
SQL> temporary tablespace temp quota unlimited on temp;
```

where:

username is the name to assign to the new user account.

password is the password to assign to the new user account.

unlimited is size of the tablespace (shown here as set to unlimited).

4. To grant the role with the necessary privileges to the user, at the command prompt, type:

```
grant asi_role to username;
```

```
conn sys as sysdba;
```

```
GRANT SELECT ON SYS.V_$LOCKED_OBJECT to username;
```

```
commit;
```

In this case, you grant `SELECT` permission to the user you created in step 3. The Oracle database server does not allow you to grant the permission to the *asi_role*. BEA WebLogic Enterprise Security uses this dynamic view to check whether one of its tables is currently being accessed. Therefore, the `SELECT` permission is required.

5. Exit SQLplus.

Using the Database Administration Utilities with Oracle

Table 3-5 lists and describes the batch and shell files provided for database administration. The files are located in the following directory:

`bea\wles42-admin\bin\`

where:

`bea` is the `BEA_HOME` directory.

`wles42-admin` is the installation directory for the Administration Application.

Table 3-5 Oracle Database Administration Utilities

File Name	Used to:
<code>export_policy_dbtype.bat</code> <code>export_policy_dbtype.sh</code>	Exports policy data. See the <i>BEA WebLogic Enterprise Security Policy Managers Guide</i> for information on how to export policy. The <code>dbtype</code> is the type of database, Sybase or Oracle.
<code>install_schema_dbtype.bat</code> <code>install_schema_dbtype.sh</code>	Installs the policy database schema. See “Installing the Policy Database Schema” on page 5-2 for information on how to install the database schema.
<code>install_sort_<dbtype>.bat</code> <code>install_sort_<dbtype>.sh</code>	Switches the sort order. When using Administration Console, the list of usernames and other policy elements can be sorted in alphabetical order or in discretionary order. This script is used to switch such sorting order. Alphabetical sort order has better performance than discretionary sort order. The parameters for this script are same as the <code>install_schema</code> script, except the parameter for sorting type, which can take value of either A (ASCII) or D (Dictionary).
<code>refresh_schema_dbtype.bat</code> <code>refresh_schema_dbtype.sh</code>	Clean up the policy created in the policy database and return it to the same state as it was following the schema installation. The parameters for this script are the same as the <code>install_schema</code> script.
<code>uninstall_schema_dbtype.bat</code> <code>uninstall_schema_dbtype.sh</code>	Uninstall the policy database schema from the database server. The parameters for this script are the same as the <code>install_schema</code> script.

Before running these scripts with an Oracle database, you need to ensure the following setup steps are completed:

- The current path (.) is in your `PATH` environment.
- Ensure that the Oracle client is set up and configured as described.
- For Windows, ensure that the `PATH` includes the `BIN` and `DLL` directory of Oracle installation.
- For Solaris, ensure that the environmental variable `ORACLE_HOME` is set, `$ORACLE_HOME/bin` is in the `PATH`, and `$ORACLE_HOME/lib` in the `LD_LIBRARY_PATH`.
- Ensure that you can connect to the Oracle database server using command `sqlplus` (the Net Service Name, login ID and password).

Backing Up an Oracle Database

BEA strongly recommends that you backup your original policy database regularly. A database backup is always recommended before you uninstall or re-install the policy database. You may need to contact your database or system administrator to assist with this process. Backups should be done on a regularly scheduled basis.

For instructions on backing up your Oracle database, see the *Oracle Backup and Recovery Guide* that comes with your Oracle documentation.

Setting Up and Administering the Sybase Database and Client

This section contains the procedures for setting up and administering an Sybase database and a Sybase Client. It covers the following topics:

- [“Before you Begin the Sybase Database Setup”](#) on page 3-32
- [“Installing and Configuring the Sybase Adaptive Server”](#) on page 3-35
- [“Installing and Configuring a Sybase Database Client”](#) on page 3-42
- [“Tuning the Sybase Database”](#) on page 3-46
- [“Administering the Sybase Policy Database”](#) on page 3-51

Before you Begin the Sybase Database Setup

Before you begin to set up your Sybase database, review the following topics to better understand Sybase database configuration requirements:

- [“Overview of the Sybase Client/Server Architecture”](#) on page 3-32
- [“Sybase Database System Requirements”](#) on page 3-34

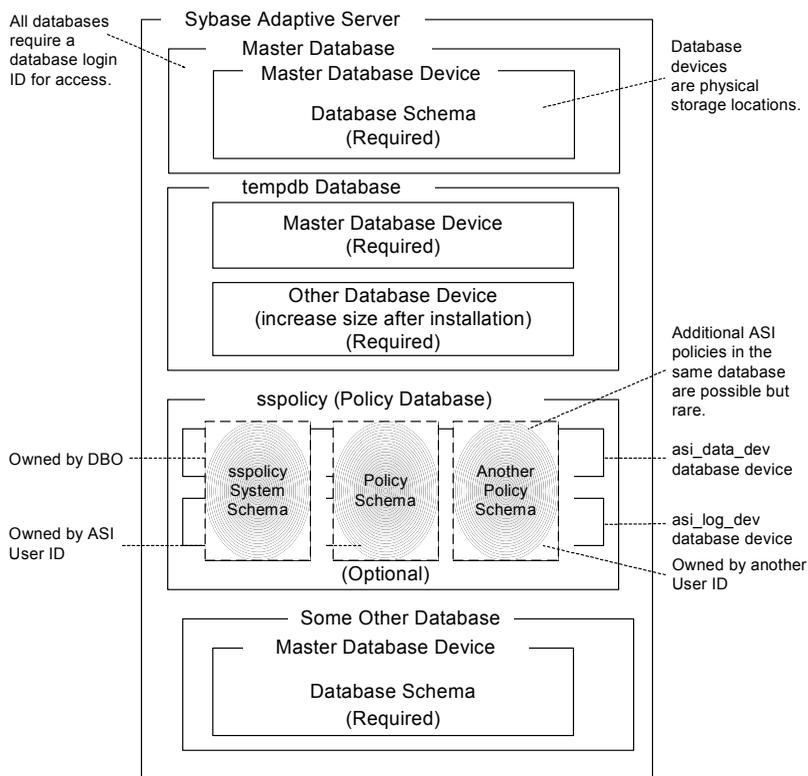
Overview of the Sybase Client/Server Architecture

The Sybase Adaptive Server is the server in the Sybase client/server architecture (see [Figure 3-5](#)). It manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory.

The policy database is a set of database schemas in which all data are stored. The Sybase database contains a set of related data tables and other database objects organized and presented to serve a specific purpose.

A database device is a Sybase term that represents the portion of a device (a portion of a hard drive, such as a partition) that is dedicated to holding database data. When creating the database device, you can choose either a raw partition or an existing file system. Choosing a raw partition can increase the performance of the database server.

Figure 3-5 Sybase Adaptive Server Setup



The Database Login ID is a login created by a system administrator to log onto the Adaptive Server. Each Database Login has a password and a default database to access. A login is valid if the Adaptive Server has an entry for that user in the system table `syslogins`.

The Database Administrator (DBA) has a special database login ID that can access all databases in the Adaptive Server. The DBA is also referred to as the system administrator. In fact, the name of the DBA login is `sa` (for System Administrator).

The Database Owner (DBO) is a special database login with permission to perform all actions on a policy database. Usually, the login that creates the database automatically becomes the DBO. The Database User ID is `dbo` (lowercase), which is different from its Database Login ID. For your policy database, you can use any Database Login ID as the DBO.

The Database User ID pertains to one specific database and is a login given permission by the DBO or DBA (system administrator) to access that one database. In most cases, the database user ID is the same as the Database Login ID. However, in some cases, they may be different, as with the special dbo user ID.

A database schema is a collection of objects associated with a particular schema name. The objects include tables, views, domains, constraints, assertions, privileges, and so on.

The policy owner is a Database User ID that controls the set of database schema in the database. BEA recommends that you not use dbo as a policy owner because it requires special administration. The WebLogic Enterprise Security architecture allows multiple policy owners in its database, each owning a policy different from the other policies.

Sybase Database System Requirements

[Table 3-6](#) describes the minimum requirements for the system on which the Sybase Adaptive Server is installed.

Table 3-6 Sybase Database Minimum Requirements

Requirement	Description
Software Version	Sybase Adaptive Server Enterprise 12.5.
Server Platform	Any platform supported by Sybase.
Memory	As required by Sybase server installation (42 MB minimum).
Disk Space for the default database	As required by Sybase server installation.
Disk Space for Sybase software	Refer to the <i>Sybase Adaptive Server Enterprise Installation Guide</i> for details.
Disk Space for the Policy Database	A minimum of two database devices is required, each having 250 MB. To approximate space requirements for any policy size, use the formula in “Calculating Sybase Database Size Requirements” on page 3-46.
Sybase Client	Sybase client that ships with Version 12.5 of the product.

Installing and Configuring the Sybase Adaptive Server

This section provides instructions for installing and configuring a Sybase database for use with the WebLogic Enterprise Security Administration Application.

For guidance on installing and configuring the database, see the following topics:

- [“Installing the Sybase Database” on page 3-35](#)—Use this procedure only if you have to install and configure the Sybase database software.
- [“Setting the Sybase Database Configuration Parameters” on page 3-38](#)—Use this procedure if the Sybase database software is already installed and you only need to set the configuration parameters as required by WebLogic Enterprise Security policy database.
- [“Creating Sybase Database Devices” on page 3-39](#)—Use this procedure to create Sybase database devices. The database devices must be created before you can create and configure the policy database.
- [“Creating and Configuring a Sybase Policy Database” on page 3-39](#)—Use this procedure to create and configure a policy database for use by the WebLogic Enterprise Security Administration Application.

Installing the Sybase Database

This section provides recommendations for installing and configuring the Sybase database software. If the Sybase database is already installed on the database host machine, you can skip this procedure and go to [“Setting the Sybase Database Configuration Parameters” on page 3-38](#).

To install the Sybase Adaptive Server, perform these steps:

1. To install a Sybase Adaptive Server database software, follow the Sybase installation instructions in the *Sybase Adaptive Server Enterprise Installation Guide*. When you are finished, go to step 2.

Note: In Sybase 12.5, you can choose the logical page size of 2K, 4K, 8K and 16K when building the server. This choice can affect the maximum length of usernames, resource names, and the length of rules, etc., when administering the security policy. See the Sybase Adaptive Server documentation for more information regarding the logical page size and column size limit.
2. Use `isql` or the Sybase Adaptive Server tool to set the `sa` password and the Sybase database configuration parameters. Do one of the following:

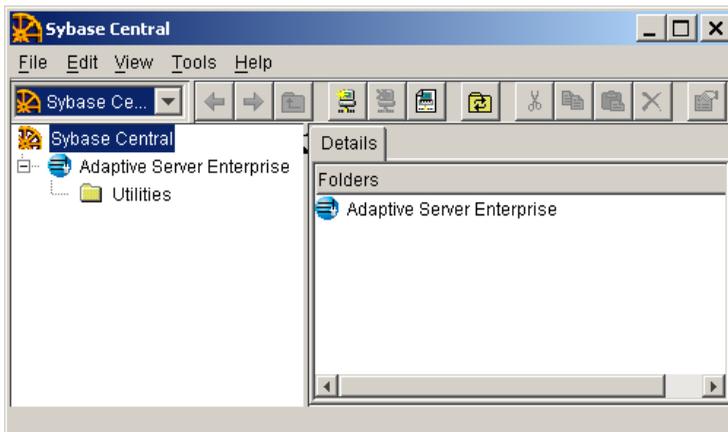
- To use `isql`, open a command console and log into the server as user `sa` and set the Sybase database configuration parameters. For instructions on how to use `isql`, see *Sybase Adaptive Server Enterprise Installation Guide*.
- To use Sybase Central Java Edition tool to set the Sybase database configuration parameters, go to the next step.

To make modifications to the server configuration, you must login as a Sybase system administrator. After making changes, you must restart the Sybase database server for the change to take effect. [Table 3-7](#) describes the settings that BEA recommends for the Sybase configuration parameters. These settings are case sensitive.

You can access Sybase Adaptive Server from the same machine as the Adaptive Server or from another client machine. To access it from a client machine, you must install the Sybase Open Client on the client machine and configure the client machine to connect to the Sybase database server (see [“Installing and Configuring a Sybase Database Client”](#) on [page 3-42](#)).

3. Click Start>Programs>Sybase>Sybase Central Java Edition. The Sybase Central control page appears ([Figure 3-6](#)).

Figure 3-6 Sybase Central Java Edition Tool



4. From the Menu Bar, click Tools and then select Connect.
5. Enter the username `sa` and click Ok. The Sybase database server appears as a node in the left pane.

6. Expand the server node, click on the Logins folder, right click on sa in the right pane and select properties. The Login Properties window appears.
7. From the Login Properties windows, select Parameters, click on Change Password, and then set the password as desired.
8. In the left pane, right click on the server node and select Configure from the drop-down menu. The Server Properties window appears.
9. Refer to [Table 3-7](#) and set the configuration parameters as directed.

Table 3-7 Sybase Configuration Parameters

Parameter	Description
Max online engines	Sets the number of processors on the host machine. The installation does not detect the number of processors on the host, so if your host has more than one (the default), you should change this parameter to reflect that. This option can increase server performance.
Max memory	Increasing this value can dramatically increase the performance of the Sybase server and Administration Console. Consult your system administrator to determine the amount of RAM available for the Sybase server.
Procedure cache size	Increase this from its default value of 25% to potentially increase server performance.
Identity burning set factor	Use a smaller value than the default value of 5000. BEA recommends a value of 1.
Number of user connections	The default is 25 connections. If your server is shared with other databases other than the policy database, you may want to increase this number. The Policy Distributor uses two connections. The server uses a pool connection. You can adjust the size of the pool connection in the configuration file. Under normal conditions, it does not use a large number of connections.

Table 3-7 Sybase Configuration Parameters (Continued)

Parameter	Description
Number of locks	<p>Increase the default value of 5000 to a larger setting if you anticipate distribution of large policies. Consider a policy of more than 2000 users and roles as a large policy. In practice, this value may be set to 10,000 or 20,000, together with the lock promotion mechanism installed later.</p> <p>Locks use a lot of memory. The Administration Console solves this problem by including a mechanism called lock promotion. If you install lock promotion, you can greatly reduce the number of locks used for distribution.</p>
Number of open indexes	Increase this number from its default setting to a value like 2000.

Setting the Sybase Database Configuration Parameters

Use this procedure this procedure to set the Sybase database configuration parameters.

Note: If you installed the Sybase database software and set these parameters as described in “[Installing the Sybase Database](#)” on page 3-35, skip this procedure and go to “[Creating Sybase Database Devices](#)” on page 3-39.

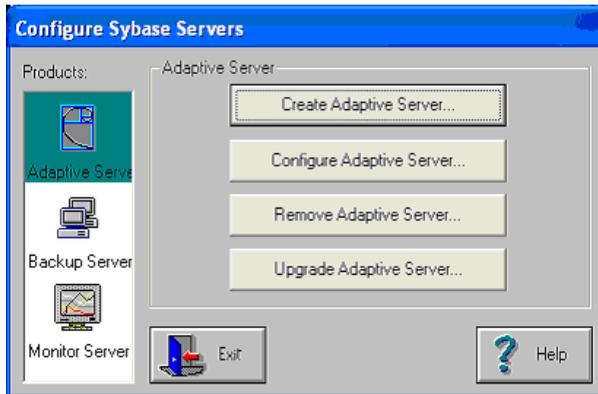
To set the Sybase database configuration parameters, perform the following steps:

Note: The section provides guidance to assist you, but it does not supersede the documentation provided by Sybase.

1. Use `isql` or the Sybase Server Config tool to set the `sa` password and the Sybase database configuration parameters. Do one of the following:
 - To use `isql`, open a command console and log into the server as user `sa` and modify the server configuration. For instructions on how to use `isql`, see *Sybase Adaptive Server Enterprise Installation Guide*.
 - To use Sybase Server Config tool to set the Sybase database configuration parameters, go to the next step.

Note: After you make the configuration changes, reboot the database server machine to have changes take effect.

2. Click Start>Programs>Sybase>Server Config. The Configure Sybase Server screen appears (see [Figure 3-7](#)).

Figure 3-7 Configure Sybase Servers Screen

3. Select the Adaptive Server product, login as user `sa`, select the Command Line Change Option, and set the Sybase configuration parameters listed and described in [Table 3-7](#)

Creating Sybase Database Devices

The policy database requires at least two database devices, each having at least 250 MB of free space. The first device stores policy data and the other stores the transaction log. You must create these two database devices before you create and configure the policy database.

Note: For better performance, BEA recommends a raw partition as the best configuration for the database device. Obviously, you must allocate sufficient disk space to ensure that the database meets your performance requirements.

For instructions on how to create Sybase database devices, see the Chapter “Managing Adaptive Server Databases” in the *Sybase Adaptive Server Enterprise Configuration Guide* for the platform on which you installed the database server: Microsoft Windows, Solaris, or Linux.

Creating and Configuring a Sybase Policy Database

Like other Sybase databases, the policy database contains at least one set of database schemas, owned by a user referred to as the policy owner. While it is unusual, the same database may contain multiple sets of policies, each owned by a different user.

Note: Before continuing, be sure that you have the names of two existing database devices that have sufficient free space to hold the data and transaction log for the policy database. If the database devices do not exist, go to “[Creating Sybase Database Devices](#)” on [page 3-39](#) and create them.

To create and configure the policy database, perform these steps:

1. From a command prompt, log into the database server as the Sybase system administrator. For example, type:

```
isql -Usa -SASERVER
```

where: *ASERVER* is the name of your database server.

2. Enter the following commands:

```
1>use master
2>go
1>create database sspolicy on asi_data_dev = 250 log on asi_log_dev =
  250
```

where: *sspolicy* is the name of the database. The name *sspolicy* is used only for the purpose of the example. You can assign any name to the database. In this example, the minimum database sizes, 250 MB, are used. If you choose to use other sizes, enter those sizes instead.

asi_data_dev and *asi_log_dev* are the names of the two devices.

```
2>go
```

3. To use the Sybase *sp_dboption* system procedure to set the database options, type the following commands at the *isql* command prompt:

```
1>use master
2>go
1>sp_dboption sspolicy, "select into/bulkcopy", true
2>go
1>sp_dboption sspolicy, "abort tran on log full", true
2>go
1>sp_dboption sspolicy, "trunc log on chkpt", true
2>go
1>sp_dboption sspolicy, "trunc. log on chkpt.", true
2>go
```

For more information on the *sp_dboption* system procedure, see *Sybase Adaptive Server Enterprise Reference Manual: Procedures*.

Note: In a development database, you may be set the *trunc log on chkpt* option to *true* because the DBA may not have time to run a dump transaction from time-to-time to truncate the transaction log. In a production database, you must set this option to *false* and perform a dump transaction to back up and truncate the database and transaction logs.

4. To create the database user account for the WebLogic Enterprise Security Administration Application to access the policy database, perform these steps:

- a. To create the ASI Database Login ID, at the `isql` command prompt, type the following commands:

```
1>use master
2>go
1>sp_addlogin asi, password, sspolicy, null, "asi login"
2>go
```

The `password` must be at least six alphanumeric characters or other characters allowed by Sybase. The name of the default database is `sspolicy`. If an `asi` login already exists, you must use the `sp_modifylogin` command to set its default database to `sspolicy`.

- b. To create the ASI Database User ID, type the following commands:

```
1>use sspolicy
2>go
1>sp_adduser asi
2>go
```

- c. To grant Permissions to the ASI Database User ID, type the following commands:

```
1>use sspolicy
2>go
1>grant all to asi
2>go
```

5. To verify that the configured user `asi` can connect to the target Sybase database using `isql`, open a command window on the machine on which the database is installed and login. For example, using the values specified in the previous step, type the following:

```
isql -Uasi -Sserver_name
Password: password
1>
```

where: `server_name` is the database server name and `password` in the password of the `asi` user.

This completes the configuration of the policy database.

Installing and Configuring a Sybase Database Client

Skip this step if you want to administer the Sybase Adaptive Server and run the WebLogic Enterprise Security Administration Application on the machine on which the Sybase Adaptive Server is installed.

You must install the Sybase Open Client (Sybase client for Adaptive Server) to:

- Administer the Adaptive Server from a machine that does not already have the Adaptive Server or Open Client installed.
- Install the database schema or run any of the servers on a machine that does not already have Adaptive Server or Open Client installed. These servers include the Administration Server (on which your administration console is running) and the Policy Distributor.

The information you need to install and configure the Sybase Open Client includes:

- Sybase Server Name
- Username and password to log in to the database server
- Hostname and port number the database server is running on
- Sybase Database name

The following topics provide guidance for installing and testing a Sybase Open Client:

- [“Testing an Existing Sybase Open Client Installation” on page 3-42](#)
- [“Installing and Configuring the Sybase Open Client on Windows” on page 3-43](#)
- [“Installing and Configuring the Sybase Open Client on Sun Solaris” on page 3-44](#)
- [“Installing and Configuring the Sybase Open Client on Red Hat Advanced Server 2.1” on page 3-45](#)

Testing an Existing Sybase Open Client Installation

If the Sybase Open Client is already installed, you need to ensure that you can access the Adaptive Server from the client. To do so, open a command window and type:

```
isql -U loginid -S ASERVER -P loginidpassword
```

where: *loginid* is the identity you defined when configured the policy database, *ASERVER* is the name of the policy database, and *loginidpassword* is the password of the identity.

The `isql` prompt appears, indicating a successful connection.

If this command fails and you know the client is installed, the client is probably not configured properly to point to the database server. If the client is on the same machine as the Sybase database, the client is configured automatically when you do the installation. If the client is on a machine other than the Sybase database machine, you need to configure the client. For instructions on how to configure the Open Client, see the installation and configuration procedure that applies to your particular platform:

- “Installing and Configuring the Sybase Open Client on Windows” on page 3-43
- “Installing and Configuring the Sybase Open Client on Sun Solaris” on page 3-44
- “Installing and Configuring the Sybase Open Client on Red Hat Advanced Server 2.1” on page 3-45

Installing and Configuring the Sybase Open Client on Windows

To install the Sybase Open Client in a Windows environment, do the following:

Note: The section provides guidance to assist you, but it does not supersede the documentation provided by Sybase.

1. Log into Windows as administrator.
2. Start the Open Client installation program on your computer (setup.exe) and install the Open Client according to instructions provided in the *Sybase Adaptive Server Enterprise Installation Guide* for Windows. If the Open Client is already installed, skip this step and go to the next step.
3. Check that your system environment variables are set correctly to point to the Sybase installation directory, as shown in the following example (where the installation is on the D: drive):

```
SYBASE=D:\Sybase
SYBASE-JRE=D:\sybase\shared-1_0\JRE-1_3
SYBASE_OCS=OCS-12_5
```

4. Check that your system `PATH` environmental variable includes the `bin` and `dll` subdirectories of your Sybase installation directory, as shown in the following example (where the installation is on the D: drive):

```
D:\Sybase\OCS-12_5\bin and D:\Sybase\OCS-12_5\dll
```

5. Using a text editor or the `Dsedit` utility provided by Sybase, edit the Sybase configuration file `sql.ini` in the `\ini` sub-folder of your Sybase Open Client installation directory to include a server entry that points to your policy database server. For instructions on how to

use the `Dsedit` utility to edit the `sql.ini` file, see the *Sybase Adaptive Server Enterprise Installation Guide* for Windows. For parameters required to edit the `sql.ini` file, see the `sql.ini` file located in `\sybase\ini` directory on the machine on which the Sybase database server is installed. Here is an example `sql.ini` file produced by the `Dsedit` utility:

```
[ASERVER]
master=TCP,PCWIZ, 5000
query=TCP,PCWIZ, 5000
```

6. To test your installation, at the command prompt, type:

```
isql -U loginid -S ASERVER -P loginidpassword
```

where: `loginid` is the identity you defined when configured the policy database, `ASERVER` is the name of the policy database, and `loginidpassword` is the password of the identity.

The `isql` prompt appears, indicating a successful connection.

This completes the configuration of the Sybase Open Client.

Installing and Configuring the Sybase Open Client on Sun Solaris

To install and configure a Sybase Open Client on Sun Solaris, perform the following steps:

Note: The section provides guidance to assist you, but it does not supersede the documentation provided by Sybase.

1. Login to Solaris with the username `sybase`. If the user `sybase` does not exist, have your Solaris system administrator create it.
2. Start the Open Client installation program to install on your workstation and install the Open Client according to instructions provided in *Sybase Adaptive Server Enterprise Installation Guide* for Solaris.
3. Set the `SYBASE` environment variable to point to the Sybase installation directory, as shown in the following example:

```
/export/home/sybase
```

4. Set the `PATH` environment variable to include the `bin` subdirectory of your Sybase installation directory, as shown in the following example:

```
/export/home/sybase/OCS-12_5/bin
```

5. Set the `LD_LIBRARY_PATH` environment variable to include the `lib` subdirectory of your Sybase installation directory, as shown in the following example:

```
/export/home/sybase/OCS-12_5/lib
```

- Using a text editor or the `Dsedit` utility provided by Sybase, edit the Sybase configuration file `sql.ini` in the `\ini` sub-folder of your Sybase Open Client installation directory to include a server entry that points to your database server. For instructions on how to use the `Dsedit` Utility to edit the `sql.ini` file, see the *Sybase Adaptive Server Enterprise Installation Guide* for Solaris. For parameters required to edit the `sql.ini` file, see the `sql.ini` file located in `\sybase\ini` directory on the machine on which the Sybase database server is installed. Here is an example `sql.ini` file produced by the `Dsedit` utility:

```
[ASERVER]
master=TCP,PCWIZ, 5000
query=TCP,PCWIZ, 5000
```

- To test your installation, at the Solaris command prompt, type:

```
isql -U loginid -S ASERVER -P loginpassword
```

where: `loginid` is the identity you defined when configured the policy database, `ASERVER` is the name of the policy database, and `loginidpassword` is the password of the identity.

The `isql` prompt appears, indicating a successful connection.

- Repeat steps 3 to 5 for each user needing access to the Sybase Adaptive Server.

Include these settings in either `.profile` or `.cshrc`, depending on the default user shell.

This completes the configuration of the Sybase Open Client.

Installing and Configuring the Sybase Open Client on Red Hat Advanced Server 2.1

To install and configure a Sybase Open Client on Red Hat Advanced Server 2.1, perform the following steps:

Note: The section provides guidance to assist you, but it does not supersede the documentation provided by Sybase.

- Install the Red Hat Advanced Server software according to instructions in the *Sybase Adaptive Server Enterprise Installation Guide*.
- To test your installation, at the command prompt, type:

```
isql -Usa -Ppassword -Sserver_name
```

where: `server_name` is the database server name and `password` in the password of the `sa` user.

The `isql` prompt appears, indicating a successful connection.

This completes the configuration of the Sybase Open Client.

Tuning the Sybase Database

After you have installed and configured the Sybase database and the Sybase Client, you should tune the database to suit the needs of your particular environment. The following topics provide information to assist in tuning your Sybase database:

- [“Calculating Sybase Database Size Requirements” on page 3-46](#)
- [“Calculating Sybase Tablespace Requirements” on page 3-47](#)
- [“Calculating Sybase Data Size Requirements” on page 3-48](#)
- [“Calculating Sybase Transaction Log Size Requirements” on page 3-49](#)
- [“Preventing Database Log Bloat with Sybase” on page 3-50](#)
- [“Expanding the Policy Database with Sybase” on page 3-50](#)
- [“Optimizing the Sybase Database for Large Policies” on page 3-50](#)

Calculating Sybase Database Size Requirements

For the policy database, allot the amount of disk space based on the size of your policy. BEA recommends 250 MB as an absolute minimum.

For the policy database transaction log, allot the size for the transaction log database by considering the following factors:

- How often the transaction log is dumped. The less frequent the dumping, the greater the size requirement.
- The greatest possible size of your distributed policy.
- How often the policy is distributed before the transaction is dumped. If the policy is distributed frequently before dumping the transaction, a larger transaction log is required.

The size of the data and transaction log can be increased later to use any database devices, by using the SQL command `alter database`.

Calculating Sybase Tablespace Requirements

To determine the tablespace size requirements, allot the amount of disk space based on the size of your policy, with a 250 MB as an absolute minimum, provided the rollback segments can handle the policy loading and distribution.

To determine your actual tablespace requirements, see the following topics:

- [“Minimum Disk Space Allotment” on page 3-22](#)
- [“Group Flattening and Rules” on page 3-47](#)
- [“Metadirectory Synchronization Services” on page 3-23](#)

Minimum Disk Space Allotment

The 250 MB minimum space works fine with a small policy and a small user community such as the following:

- The policy has a maximum of 1000 users
- Each user belongs to no more than one group
- Each user has one single-valued attribute
- The policy has less than 100 privileges, resources, and declarations
- The policy has less than 100 flattened rules; no composite privileges, resources, or subjects (users and groups) in the rule

Group Flattening and Rules

Group flattening means that a rule can exist in one of two forms: a simple rule or a composite rule. A composite rule is a combination of two or more simple rules to make them easier to use. The process for reducing a composite rule to its component simple rules is called "flattening the group."

For example, if you had three local users named Joe, Betty, and Sam, you could grant those users a role in an application by creating a composite rule like this:

```
Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,
      [//user/AcctDept/Joe/, //user/AcctDept/Sam/, //user/AcctDept/Betty/]);
```

In the policy language, this rule means "grant Joe, Sam, and Betty, who belong to the AcctDept, the role of bookkeeper in the accounting application, AcctApp."

The rule is a composite rule because it reduces or flattens to these three simple rules:

```
Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,  
//user/AcctDept/Joe/); and  
  
Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,  
//user/AcctDept/Sam/); and  
  
Grant (//role/bookkeeper, //app/policy/AcctDept/AcctApp,  
//user/AcctDept/Betty/);
```

Even though you may see one composite rule, the composite is actually stored and distributed as three flattened simple rules. The main ramification of rule flattening is that your policies can take much more disk space than you might think when simply looking at your policy. For information on how the policy rules and how to construct rules, see [Securing Resources and Defining Policy Rules](#) in the *Policy Managers Guide*.

Metadirectory Synchronization Services

If you want to use the BEA WebLogic Enterprise Security Metadirectory Synchronization Services, you must create an additional set of tables to use to synchronize identity information. The amount of space required to store identity information approximately doubles so you should allocate an appropriate amount of extra tablespace. For more information, see [Configuring Metadirectories](#) in *WebLogic Enterprise Security Administration Application Installation*.

Calculating Sybase Data Size Requirements

You can estimate your space requirements using the following formulas. With group flattening, like rules, group memberships are also reduced or flattened to their simple data components. For example, if you have a user that belongs to a group through group inheritance, the membership is stored as though the user were a direct member of the group. Thus, there is a separate group to user mapping for each group in the inheritance hierarchy. All numeric results are represented in megabytes. All formulas use the variables described in [Table 3-8](#).

Table 3-8 Sybase Variables

Variable	Description
a	Total number of user attribute values for all users, in thousands
d	Total number of declarations, in thousands
m	Total number of flattened user/group mappings, in thousands
o	Total number of objects, in thousands
p	Total number of privileges, in thousands

Table 3-8 Sybase Variables (Continued)

Variable	Description
q	Total number of object attribute values for all resources, in thousands
r	Total number of flattened rules, in thousands
u	Total number of users, in thousands

Use the following formula to calculate your data size requirements. For a description of the formula variables, see [Table 3-8](#).

$$\text{Disk Space} = 250 + 0.3u + 0.2a + 0.1m + 1.2(o + p) + 0.75(q-1) + 4d + 5r$$

For example, if all the variables had the value 5, the formula looks like this:

$$= 250 + 0.3(5) + 0.2(5) + 0.1(5) + 1.2(5 + 5) + 0.75(5-1) + 4(5) + 5(5)$$

and reduces to this:

$$= 250 + 1.5 + 1 + 0.5 + 12 + 3 + 20 + 25$$

and finally:

$$= 313$$

Thus, the example requires a minimum of 313 MB of disk space.

Note: If your server has logical page size other than 2K, increase this space proportionately.

Calculating Sybase Transaction Log Size Requirements

Use the following formula to calculate log size requirements. For a description of the formula variables, see [Table 3-8](#).

$$\text{Disk Space} = 250 + 2.5(o + p) + 2.5(q-1) + 6d + 10r$$

This formula represents the size needed for loading and distribution at once before dumping the transaction log. Once the log is dumped after loading, the space requirement drops by a third.

Note: Contact your Database Administrator to find out the actual database device usage and for assistance on extending the device size or to adding a device. If your server has a logical page size other than 2K, you need to increase this space in proportion.

Preventing Database Log Bloat with Sybase

BEA recommends that you regularly backup your policy databases. If you fail to do so, the transaction log can become quite large and could become so full that the database stops functioning. If you set the `trunc log on chkpt` database option to true, you will not have to manually dump the log from time to time. If you do want to manually dump the database or transaction logs, use the `dump database` and `dump transaction` commands. See your *Sybase Administration Guide* for more information.

Expanding the Policy Database with Sybase

If your policy grows, you may need to expand your policy database. To do so, use the `alter database` command. If there is no more free space on any of your Sybase database devices, you may need to create a new device. To do so, use the `disk init` command.

If you do create a new database device, be sure not to combine the data and log databases on the same database device. See your *Sybase SQL Server Reference Manual* for more information.

Optimizing the Sybase Database for Large Policies

When your database must contain a large policy, you may want to do one or more of the following to optimize performance:

- If your server has multiple processors, ensure that the max online engines setting reflects the number of processors.
- Ensure that you allot the maximum amount of RAM to the Sybase server.
- Ensure that you allot enough disk space for the data and transaction logs.
- Increase your `tempdb` size to facilitate sorting of large data sets.
- Run `lockpromotion_sybase.bat` or `lockpromotion_sybase.sh` to install the lock promotion mechanism to facilitate policy distribution.
- Regularly backup your database.
- Regularly dump the transaction log.
- Run `install_sort_sybase.bat` or `install_sort_sybase.sh` to enable ASCII sorting, instead of the dictionary sorting that comes with the default database schema installation. This improves the Administration Console response time.

Administering the Sybase Policy Database

This section covers the following database administration topics:

- [“Creating a User Account in a Sybase Policy Database” on page 3-52](#)
- [“Using the Database Administration Utilities with Sybase” on page 3-53](#)
- [“Backing Up a Sybase Database” on page 3-55](#)

Creating a User Account in a Sybase Policy Database

This section describes how to configure a new user account in a Sybase database. This account is necessary so that the policy for the instance of the Administration Application managed by this user can have a dedicated storage area allocated in the database instance.

To set up the user account, create the login to the Adaptive Server Enterprise database, create the user for policy database, and grant the user privileges to manipulate the policy schema.

Note: BEA strongly recommends that you not use the `dbo` of the policy database as the policy owner. While it is possible to do so, it requires additional database configuration that is beyond the scope of this guide.

To create a database user account, perform these steps:

1. Log in as the System Administrator.

2. At the command prompt, type:

```
isql -Usa -S server_name
```

where: *server_name* is the database server name.

3. To create the ASI Database Login ID, at the `isql` command prompt, type the following commands:

```
1>use master
2>go
1>sp_addlogin asi, password, sspolicy, null, "asi login"
2>go
```

where: *password* must be at least six alphanumeric characters or other characters allowed by Sybase and *sspolicy* is the name of the default database. If an `asi` login already exists, you must use the `sp_modifylogin` command to set its default database to *sspolicy*.

4. To create the ASI database user ID, at the `isql` command prompt, type the following commands:

```
1>use sspolicy
2>go
1>sp_adduser asi
2>go
```

5. To grant permissions to the ASI database user ID, at the `isql` command prompt, type the following commands:

```
1>use sspolicy
2>go
```

```
1>grant all to asi
2>go
```

Using the Database Administration Utilities with Sybase

Table 3-9 lists and describes the batch and shell files provided for database administration. The files are located in the following directory:

```
bea\wles42-admin\bin\
```

where:

bea is the BEA_HOME directory.

wles42-admin is the installation directory for the Administration Application.

Table 3-9 Database Administration Utilities

File Name	Used to:
export_policy_dbtype.bat export_policy_dbtype.sh	Exports policy data. See the <i>BEA WebLogic Enterprise Security Policy Managers Guide</i> for information on how to export policy. The <i>dbtype</i> is the type of database, Sybase or Oracle.
install_schema_dbtype.bat install_schema_dbtype.sh	Installs the policy database schema. See “ Installing the Policy Database Schema ” on page 5-2 for information on how to install the database schema.
install_sort_dbtype.bat install_sort_dbtype.sh	Switches the sort order. When using Administration Console, the list of usernames and other policy elements can be sorted in alphabetical order or in discretionary order. This script is used to switch such sorting order. Alphabetical sort order has better performance than discretionary sort order. The parameters for this script are same as the <code>install_schema</code> script, except the parameter for sorting type, which can take value of either A (ASCII) or D (Dictionary).
refresh_schema_dbtype.bat refresh_schema_dbtype.sh	Clean up the policy created in the policy database and return it to the same state as it was following the schema installation. The parameters for this script are the same as the <code>install_schema</code> script.
uninstall_schema_dbtype.bat uninstall_schema_dbtype.sh	Uninstall the policy database schema from the database server. The parameters for this script are the same as the <code>install_schema</code> script.

Table 3-9 Database Administration Utilities (Continued)

File Name	Used to:
lockpromotion_sybase.bat lockpromotion_sybase.sh	Install the lock promotion mechanism to facilitate distribution of large policy in Sybase. See “Expanding the Policy Database with Sybase” on page 3-50 for details. You need DBA access to the database to run this script.
unlockpromotion_sybase.bat unlockpromotion_sybase.sh	Uninstall the lock promotion mechanism performed by lockpromotion_sybase. See “Expanding the Policy Database with Sybase” on page 3-50 for details. You need DBA access to the database to run this script.

Before running these scripts with a Sybase database, you need to ensure the following setup steps are completed:

- The current path (.) is in your PATH environment.
- Ensure Sybase 12.5 client is set up and configured as in the Database Setup.
- Ensure that the SYBASE environmental variable is set.
- In Windows, ensure that PATH includes %SYBASE%\OCS-12_5\bin and %SYBASE%\OCS-12_5\dll.
- In Solaris, ensure that PATH includes \$SYBASE/OCS-12_5/bin and that LD_LIBRARY_PATH includes \$SYBASE/OCS-12_5/lib.
- Ensure you can connect to the Sybase database server using the isql command (the name of the database server, login ID and password).

Backing Up a Sybase Database

BEA strongly recommends that you backup your original policy database regularly. A database backup is always recommended before you uninstall or re-install the policy database. You may need to contact your database or system administrator to assist with this process. Backups should be done on a regularly scheduled basis.

If you have an existing backup procedure in place, you may choose to run it. Otherwise, follow these steps:

1. Login to your Sybase database server as the system administrator, database operator, or database owner.

The database owner is not the same as the policy owner.

2. Backup the transaction log by using the Sybase `dump transaction` command.
3. Backup the database by using the Sybase `dump database` command.

Note: See your Sybase documentation for further information on using these commands.

Setting Up and Administering the Database

Installing

The following sections describe how to install the Administration Application on both Windows and UNIX systems:

- “Before you Begin” on page 4-1
- “Starting the Installation Program on Windows Platforms” on page 4-7
- “Starting the Installation Program on a Sun Solaris Platform” on page 4-8
- “Starting the Installation Program on a Linux Platform” on page 4-11
- “Running the Installation Program” on page 4-12
- “Installing a Secondary Administration Application” on page 4-19

Before you Begin

Before you begin this installation procedure, make sure to do the following:

- Download and read the Release Notes
<http://e-docs.bea.com/wles/docs42/relnotes/download.html>
- Ensure that your computer meets all the prerequisites described in “Preparing to Install” on page 2-1.
- Install WebLogic Server 8.1, with Service Pack 3 or Service Pack 4
- Install and Configure the Database (Oracle or Sybase)
- Install the Database Client Software
- Obtain secure user names and passwords

The following topics provide additional information to assist you in preparing for an installation:

- [“System Security and BEA Weblogic Enterprise Security” on page 4-2](#)
- [“Secure User Names and Passwords” on page 4-3](#)
- [“Generating a Verbose Installation Log” on page 4-6](#)

System Security and BEA Weblogic Enterprise Security

Like any component running on a system, the infrastructure it provides is only as secure as the operating environment where it is installed. When BEA WebLogic Enterprise Security is installed on a system, it makes use of that system's security infrastructure to lock itself down and integrate with the security of its environment. Through the use of user, group, and file system permissions, BEA WebLogic Enterprise Security allows limited access to many operations depending upon these permissions. For more information on users, groups, and file system permission, see the following topics:

- [“System Users” on page 4-2](#)
- [“System Groups” on page 4-2](#)
- [“File System Permissions” on page 4-3](#)

System Users

BEA WebLogic Enterprise Security uses two user identities when installed on a system. These identities are selected when the first product in BEA WebLogic Enterprise Security family is installed and are referred to as the Administration User and the Service Control Manager User.

The Service Control Manager user is the identity assumed by the Service Control Manager when it starts. The Service Control Manager is the component that brokers trust between the local system and the Administration Server.

The other identity on the system is the Administration User. The Administration user owns all files (other than the Service Control Manager) and, on an Administration Server, is the identity the Administration Server assumes when it starts.

System Groups

Two groups are used in addition to the user identities to secure the Application Security Infrastructure.

- The Security Administrators Group allows users other than the Service Control Manager user or the Administrative user to perform log maintenance, creation and destruction of new instances, and other administrative tasks.
- The Security Users Group permits users on the system to have the necessary permissions to use and execute applications protected by WebLogic Enterprise Security. All WebLogic Enterprise Security users, including administrators, must belong to this group.

File System Permissions

File system permissions are used to enforce user and group based restrictions. With each product, and instance a lockdown script is created and run when installation occurs: `lockdown.bat` (Windows) or `lockdown.sh` (Unix or Linux). This lockdown script can be run again at a later time to restore the installation to the recommended file system permissions.

There are two directories that contain executable tools and utilities: `adm` and `bin`. The `adm` directory contains tools and utilities that only an administrator can run, for example `enrollment`. The `bin` directory contains tools and utilities that all security users can run, for example `set-env`. The `log` directory is writable by all security users, but can only be read by security administrators (or on UNIX, only the instance owner). The `work` directory is a temporary directory that can only be read and written to by security users.

Secure User Names and Passwords

WebLogic Enterprise Security implements a sophisticated username and password schema to protect the application itself and to ensure secure communications. Understanding this schema is important to installing the product and ensuring that it operates properly in either a development or production environment.

There are three levels of password protection: local system usernames and passwords (protect the WebLogic Enterprise Security components), passwords for keystores (secure communication between components), and a password to protect the private keys (the Certificate Authority). Understanding your enterprise and how responsibilities in your the organization are separated is essential to establishing a secure environment. For example, the person who maintains the database is usually not the person who designs and implements security. The person who deploys applications is usually not the person who administers system usernames and passwords. And, while you may not be as concerned with a more formal authorization in your development environment, your production environment needs to be firmly secured and responsibilities clearly defined.

WebLogic Enterprise Security user accounts on Windows platforms, like `asiadmin` and `scmuser` are special (see [System Users](#) and [System Groups](#)), and cannot be used to logon to any interactive session; these passwords are used for registration purposes only. They can only be used to start and stop component services. After the installer collects all of the passwords, it encrypts them in an internal password file. Later on, the service engine uses the username and password to register WebLogic Enterprise Security as a Windows service. So, the user may not need to change the password for the newly created specific user names like `asiadmin` and `scmuser`; but, optionally, they can be changed if necessary.

Note: If these user names already exist (they were generated as a part of a previous install process), you must enter the correct password. Remember to write down all usernames and passwords and store them in a safe place.

User names and passwords are required to access the following components. [Table 4-1](#) lists each component that requires a value and lists the default values.

Table 4-1 User Names and Passwords

Component	Description	Default
Database Server	A database server account used to connect to the database server where the policy data is stored, and update policy data using the policy import and export tools.	none
Administration Server	A local user account used to start the Administration Server and all Administration Application components.	User: <code>asiadmin</code> Group: <code>asiadgrp</code>
Service Control Manager	A local user account used to start the Service Control Manager.	User: <code>scmuser</code>
Security Group	A local group that includes all users of WebLogic Enterprise Security. All users of WebLogic Enterprise Security must be members of this security group, including administrators.	Group: <code>asiusers</code>
Certificate Authority	Sets the password for the private key for the Certificate Authority. All trust within the enterprise domain originates from this authority.	Randomly generated

Table 4-1 User Names and Passwords (Continued)

Component	Description	Default
Identity Key Passwords (Keystore Passwords)	<p>You also need to supply private key passwords for each of the following identities:</p> <ul style="list-style-type: none"> • Service Control Manager • Security Service Module • Administration Application <p>Private key passwords validate process authenticity by using the Certificate Authority chain of trust. Identities with invalid or untrusted keys cannot participate in the trust relationships in the enterprise domain.</p>	Randomly generated
Configure Keystores	<p>You need to supply keystore passwords for each of the Identity, Peer and Trust keystores.</p> <p>Identity Keystore - stores and protects the private keys that represent the processes identity or identities.</p> <p>Peer Keystore - stores and protects the public keys for all trusted identities within the installed component (Administration Application, Security Service Module or Service Control Manager).</p> <p>Trust Keystore - stores and protects public keys for Certificate Authorities that originate the chain of trust.</p>	Randomly generated

BEA recommends following these guidelines:

- **Development Environment**—In a development environment, you can either use the default values generated during the installation process or you can assign your own user names and passwords to protect your public and private keys.
- **Production Environment**—In a production environment, you must choose all passwords explicitly. These passwords may be needed for future maintenance of the public key infrastructure (PKI), for example, in the case of a failure. Make sure to write down all password information and retain it in a secure location.

Note: BEA does not recommend the use of randomly generated passwords, as the generation mechanism for these passwords is *not* secure. In a production environment, BEA does not recommend installing Security Service Modules on the same machine as the Administration Server.

Generating a Verbose Installation Log

If you start the installation process from the command line or from a script, you can specify the `-log` option to generate a verbose installation log. The installation log lists messages about events that occur during the installation process, including informational, warning, error, and fatal messages. This can be especially useful for silent installations.

Note: You may see some warning messages in the installation log. However, unless there is a fatal error, the installation program completes the installation successfully. The installation user interface indicates the success or failure of the installation, and the installation log file includes an entry indicating that the installation was successful.

To create a verbose log file during installation, use the following command line or script:

- For Windows platforms:

```
wles422admin_win32.exe -log=D:\bea\logs\wles_install.log
```

- For the Sun Solaris platform:

```
wles422admin_solaris32.bin -log=/bea/logs/wles_install.log
```

- For the Linux Red Hat Advanced Server platform:

```
wles422admin_linux32.bin -log=/bea/logs/wles_install.log
```

Note: The `-log` parameter is optional. By default, the installation log is put in the log directory where you install the Administration Server.

The path must be the full path to a file name. If the file does not exist, all folders in the path must exist before you execute the command or the installation program does not create the log file.

Starting the Installation Program on Windows Platforms

Note: Do *not* install the software from a network drive. Download the software to a local drive on your machine and install it from there.

Before running the installer, ensure the following two things are done.

- Ensure the database client directories have the correct permissions.

Set the file permissions on the database client directories so that all users can read and execute the files. Run the following command:

```
cacls C:\oracle /T /E /G Everyone:F
```

Where:

C:\oracle is the location of your database application

- Ensure the PATH is set correctly.

It is also important to include the `/bin` directory of your database client in the system `PATH` (the path available to all users) rather than the `user PATH` (the path only available to the current user). If this is changed, you must reboot before the change becomes available to processes running as services (which is how the Administration Application initializes itself).

To install the application in a Microsoft Windows environment:

1. Shut down any programs that are running.
2. Log in to the local Administrators group.
3. If you are installing from a CD-ROM, go to step 4. If you are installing by downloading from the BEA web site:
 - a. Go to <http://commerce.bea.com/showallversions.jsp?family=WLES> and download the installation file for your platform.
 - b. Go to the directory where you downloaded the installation file and double-click `wles422admin_win32.exe`.
The WebLogic Enterprise Security Administration Application window appears as shown in [Figure 4-1](#).
4. If you are installing from a CD-ROM:
 - a. Insert Disk 1 into the CD-ROM drive.

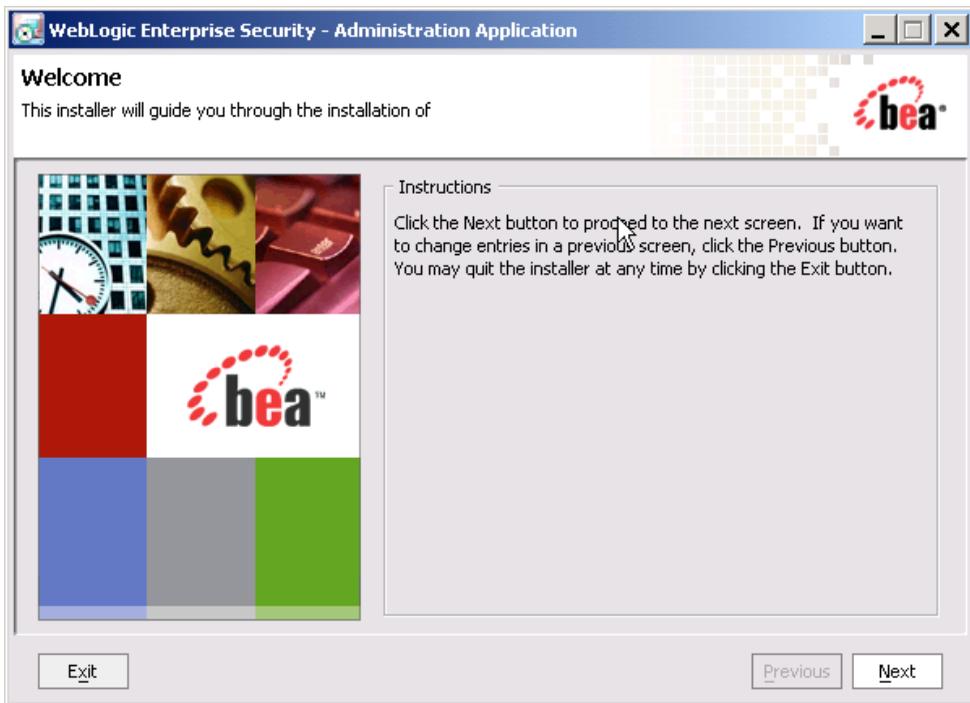
If the installation program does not start automatically, open Windows Explorer and double-click the CD-ROM icon.

- b. From the installation CD, double-click `wles422admin_win32.exe`.

The WebLogic Enterprise Security Administration Application window appears as shown in [Figure 4-1](#).

- 5. Proceed to [“Running the Installation Program”](#) on page 4-12.

Figure 4-1 WebLogic Enterprise Security Administration Application Window



Starting the Installation Program on a Sun Solaris Platform

To run graphical-mode installation, your console must support a Java-based GUI. If the installation program determines that your system cannot support a Java-based GUI, the installation program automatically starts console-mode installation.

Before running the installer, ensure the following three things are done.

- Ensure the database client directories have the correct permissions.

Set the file permissions on the database client directories so that all users can read and execute the files. Run the following command:

```
chmod -R o+rx /opt/ora92
```

- Ensure the PATH is set correctly.

It is also important to add the `/bin` directory to `PATH` and the `/lib` directory to `LD_LIBRARY_PATH`. If these settings are changed, you must reboot before the changes become available to processes running as services (which is how the Administration Application initializes itself).

Note: BEA recommends setting these variables in `/etc/profile` so they are available to all processes starting from `init`.

- Ensure that the location into which you do the install is accessible to all users at both the parent and the child directory levels.

For example, if the installation directory is `/opt/beahome/wles42-admin` and the `/opt/` directory is only accessible by root, post installation scripts that run as a user other than root cannot access the directory where they reside. Therefore, the directory into which you do the install (for example, `/opt/beahome/wles42-admin`) must have execute permissions for `other`. Run the following command to reset the permissions:

```
chmod o+x /opt/
```

The `beahome` and `wles42-admin` directories already have permissions set appropriately.

To install the application on a Sun Solaris platform:

1. Log in to the machine as root.
2. Set your `DISPLAY` variable if needed.
3. If you are installing from a CD-ROM, go to step 4. If you are installing by downloading from the BEA web site:
 - a. Go to <http://commerce.bea.com/showallversions.jsp?family=WLES> and download `wles422admin_solaris32.bin`.
 - b. Go to the directory where you downloaded the file and change the protection on the install file:


```
chmod u+x wles422admin_solaris32.bin
```
 - c. Start the installation: `./wles422admin_solaris32.bin`

Installing

The WebLogic Enterprise Security Administration Application window appears as shown in [Figure 4-1](#).

4. If you are installing from a CD-ROM:

a. Insert Disk 1 into the CD-ROM drive.

b. From the installation CD, execute `wles422admin_solaris32.bin`.

The WebLogic Enterprise Security Administration Application window appears as shown in [Figure 4-1](#).

5. Proceed to [“Running the Installation Program” on page 4-12](#).

Starting the Installation Program on a Linux Platform

To run graphical-mode installation, your console must support a Java-based GUI. If the installation program determines that your system cannot support a Java-based GUI, the installation program automatically starts console-mode installation.

Before running the installer, ensure the following three things are done.

- Ensure the database client directories have the correct permissions.

Set the file permissions on the database client directories so that all users can read and execute the files. Run the following command:

```
chmod -R o+rx /opt/ora92
```

- Ensure the PATH is set correctly.

It is also important to add the `/bin` directory to `PATH` and the `/lib` directory to `LD_LIBRARY_PATH`. If these settings are changed, you must reboot before the changes become available to processes running as services (which is how the Administration Application initializes itself).

Note: BEA recommends setting these variables in `/etc/profile` so they are available to all processes starting from `init`.

- Ensure that the location into which you do the install is accessible to all users at both the parent and the child directory levels.

For example, if the installation directory is `/opt/beahome/wles42-admin` and the `/opt/` directory is only accessible by root, post installation scripts that run as a user other than root cannot access the directory where they reside. Therefore, the directory into which you do the install (for example, `/opt/beahome/wles42-admin`) must have execute permissions for `other`. Run the following command to reset the permissions:

```
chmod o+x /opt/
```

The `beahome` and `wles42-admin` directories already have permissions set appropriately.

To install the application on a Linux platform:

1. Log in to the machine as root.
2. Set your `DISPLAY` variable if needed.
3. If you are installing from a CD-ROM, go to step 4. If you are installing by downloading from the BEA web site:

- a. Go to <http://commerce.bea.com/showallversions.jsp?family=WLES> and download `wles422admin_linux32.bin`.
- b. Go to the directory where you downloaded the file and change the protection on the file:

```
chmod u+x wles422admin_linux32.bin
```
- c. Start the installation: `./wles422admin_linux32.bin`

The WebLogic Enterprise Security Administration Application window appears as shown in [Figure 4-1](#).

4. If you are installing from a CD-ROM:
 - a. Insert Disk 1 into the CD-ROM drive.
 - b. From the installation CD, execute `wles422admin_linux32.bin`.

The WebLogic Enterprise Security Administration Application window appears as shown in [Figure 4-1](#).

5. Proceed to [“Running the Installation Program” on page 4-12](#).

Running the Installation Program

The installation program prompts you to enter specific information about your system and configuration, as described in [Table 4-2](#).

Note: You must install the Administration Application first, before installing your Security Service Modules. BEA does not recommend installing Security Service Modules on the same machine as the Administration Server in a production environment.

To complete this procedure you need the following information:

- For Windows, a username and password for the Administration Application account
- For Windows, a username and password for the Service Control Manager account
- Name of the `BEA_HOME` directory
- Name of the product directory
- Database connection information (see your database administrator and [“Setting Up and Administering the Database” on page 3-1](#) for details).

Table 4-2 Administration Application Installation

In this Window:	Perform this Action:
Welcome	Click Next to proceed or cancel the installation at any time by clicking Exit .
BEA License Agreement	Read the BEA Software License Agreement, and then select Yes to indicate your acceptance of the terms of the agreement. To continue with the installation, you must accept the terms of the license agreement, click Yes , and then click Next .
Choose BEA Home Directory	Specify the BEA Home directory that serves as the central support directory for all BEA products installed on the target system. If you already have a BEA Home directory on your system, you can select that directory (recommended) or create a new BEA Home directory. If you choose to create a new directory, the installer program automatically creates the directory for you.
Choose Product Directory	Specify the directory in which to install the Administration Application software. You can accept the default product directory (<code>wles42-admin</code>) or create a new product directory. If you choose to create a new directory, the installation program automatically creates the directory for you, if necessary. Click Next to continue.
Choose Service Control Manager Directory	Specify the directory in which to install the Service Control Manager. You can accept the default directory (<code>wles42-scm</code>) or you can create a new one. Click Next to continue.

Table 4-2 Administration Application Installation (Continued)

In this Window:	Perform this Action:
Select Users and Groups	<p>Specify the user names and group names to use for the Service Control Manager and Administration Application. You can accept the default settings or create new ones.</p> <p>Note: When installing this product for use in a production environment, BEA recommends that you set these passwords to known values; otherwise you will not be able to modify them later. For example, you may want to modify these passwords to comply with organizational requirements.</p> <p>Admin User (<i>asiadmin</i>) A local user account used to start the Administration Application components.</p> <p>Admin Group (<i>asiadgrp</i>) Administration Application group. Members of this group have full access to Administration Application and log files; they can start and stop the Administration Application components.</p> <p>SCM User (<i>scmuser</i>) A local user account used to start the Service Control Manager.</p> <p>Security Group (<i>asiusers</i>) Service Control Manager Group. Members of this group are allowed to use the WebLogic Enterprise Security product.</p> <p>Click Next to continue.</p>
Confirm User Selection	<p>If the name of the user and group do not yet exist, they are created for you. Verify the values you entered are correct, and then click Next.</p>
Select User Passwords (Windows only)	<p>Specify the password for the Administration Application User and Service Control Manager User. You can also accept the default passwords that are randomly generated.</p> <p>Note: If any of the users exist you must enter their passwords; the passwords are not generated randomly.</p> <p>Note: Passwords are case sensitive. If you are installing the Administration Application in a production environment, BEA recommends using secure user names and passwords, and not those that are randomly generated.</p> <p>Click Next to continue.</p>

Table 4-2 Administration Application Installation (Continued)

In this Window:	Perform this Action:
Choose Network Interfaces	<p>Select the network interfaces to which to bind the Service Control Manager. This is the IP Address used to listen for requests to provision policy and configuration data.</p> <p>Note: If you are installing the Administration Application in a production environment with more than one network card, you want to select a protected (internal) interface; you do not want to expose the Service Control Manager through a public address.</p> <p>Click Next to continue.</p>
Configure Administration Application	<p>Enterprise Domain Name</p> <p>Enter the name to assign to this domain. The Enterprise Domain represents the collection of Security Service Modules administered by this BEA WebLogic Enterprise Security Administration Application. Make a note of the Enterprise Domain Name you entered as you will need this to install any subsequent security service modules.</p> <p>Note: The Enterprise Domain Name <i>must</i> be entered in all lower case; and may not contain any spaces or punctuation marks.</p>
Configure Administration Application (Continued)	<p>Administration Application</p> <p>HTTP Port</p> <p>Enter the HTTP port number for the WebLogic Server 8.1 Administration Console to use.</p> <p>SSL Port</p> <p>Enter the HTTPS port number for the Administration Application to use. When you enter the SSL port number, make sure that at least six consecutive port numbers are also available. These port numbers are used by services required by the BEA WebLogic Enterprise Security Administration Application to operate properly, and the Administration Application always runs on a secure connection using this port.</p>

Table 4-2 Administration Application Installation (Continued)

In this Window:	Perform this Action:
Configure Administration Application (Continued)	<p>Certificate Authority Duration (years)</p> <p>Enter the number of years the security certificate remains in effect. The Certificate Authority is used to generate and sign certificates for other components in the BEA WebLogic Enterprise Security system.</p> <p>Secondary Server URL</p> <p>This URL is only necessary if you plan on installing the Security Service Modules on the same machine as your Administration Server and plan on configuring the Security Service Modules with a backup Administration Application. Otherwise, you can leave this URL blank.</p> <p>Click Next to continue.</p>
Configure Database Connection	<p>Database Client</p> <p>Select the type and version of database client you are using (Sybase or Oracle) on this machine. The prompts that appear differ depending on the type of client you select.</p> <p>Database Connection</p> <p>For Oracle:</p> <p>Oracle Service Name</p> <p>Local service name (Oracle System Identifier SID).</p> <p>Database JDBC URL</p> <p>Change the <SID> and <SERVER> name to complete the JDBC URL: <code>jdbc:oracle:thin:@<SERVER>:1521:<SID></code></p> <p>Database JDBC Driver</p> <p>The Oracle driver to use by default: <code>oracle.jdbc.driver.OracleDriver</code></p>

Table 4-2 Administration Application Installation (Continued)

In this Window:	Perform this Action:
Configure Database Connection (Continued)	<p>Database Connection</p> <p>For Sybase:</p> <p>Sybase Host Name Sybase server entry you configured in this local machine, used to connect to Sybase database server running elsewhere.</p> <p>Sybase Database Name Name of the Sybase database; that is, the name of policy database.</p> <p>Database JDBC URL Change the <hostname_or_IP> and <databaseName> name to complete the JDBC URL, assuming the Sybase server is running on port 4100: The <hostname_or_IP> is the hostname or IP address of the machine running Sybase server, and <databaseName> is the policy database name. You may need to change port number if necessary. The Sybase server usually listens on port 5000 on Windows platform and 4100 on other platforms: jdbc:sybase:Tds:<hostname_or_IP>:4100/<databaseName> or jdbc:sybase:Tds:<hostname_or_IP>:4100 You can use the later when the default database for Login ID is set to the policy database.</p> <p>Database JDBC Driver The Sybase driver to use by default: com.sybase.jdbc2.jdbc.SybDriver</p>
Configure Database Connection (Continued)	<p>Database Login</p> <p>Login ID, Password, Confirm Password The database login id and password to use to connect to the database; you must confirm the password.</p> <p>Click Next to continue.</p>

Table 4-2 Administration Application Installation (Continued)

In this Window:	Perform this Action:
Configure Certificate Authority	<p>The Certificate Authority is used to generate and sign certificates for other components in the BEA WebLogic Enterprise Security system.</p> <p>Key Password</p> <p>You can either choose to use a randomly generated password or you can specify the private key password. You must confirm the password.</p> <p>Note: You should write down or remember all passwords and store them in a safe location should you ever need to use them again. For example, if you plan on installing redundant servers, you need to use the same keystore and key passwords.</p> <p>Click Next to continue.</p>
Configure Keys	<p>Enter the following key passwords to secure communications of internal processes. These are components of the Administration Application. Private key passwords are used to validate process authenticity by using the Certificate Authority chain of trust. Identities with invalid or untrusted keys cannot participate in the trust relationships of the enterprise domain.</p> <ul style="list-style-type: none"> • Service Control Manager • Security Service Module • Administration Application

Table 4-2 Administration Application Installation (Continued)

In this Window:	Perform this Action:
Configure Keystores	<p>You need to supply keystore passwords for each of the Identity, Peer and Trust keystores.</p> <p>Identity Keystore stores and protects the private keys that represent the processes identity or identities.</p> <p>Peer Keystore stores and protects the public keys for all trusted identities within the installed component (Administration Application, Security Service Module or Service Control Manager).</p> <p>Trust Keystore stores and protects public keys for Certificate Authorities that originate the chain of trust.</p>
Installation Complete	<p>This page indicates the Administration Application completed successfully. If you want to install the database schema now.</p> <p>Check the Install Database Schema check box. If you are installing a failover server for backup and failover purposes, you do not want to install the database schema again, so clear this check box.</p> <p>Note: Make sure to write down the name of the Administration Application URL. You will need this URL when you are installing additional components.</p> <p>Click Done to complete the installation.</p>

What's Next

Now that you have installed the necessary software, you must start the necessary services. For additional instructions, see [“Post Installation Tasks” on page 5-1](#). If you want to install a second Administration Server to use as a backup, see [“Installing a Secondary Administration Application” on page 4-19](#).

Installing a Secondary Administration Application

You may want to install and configure a second Administration Application to support failover. For information on failover considerations, see [“Failover and System Reliability”](#) in the *BEA WebLogic Enterprise Security Administration Guide*.

The secondary Administration Server must be set up in the same manner as the primary and installed on a separate machine.

1. Install the Weblogic Server 8.1.
2. Run the installation program on the secondary Administration Server.
3. Enter the following information:
 - a. When prompted for the Enterprise Domain, make sure to enter the same domain name that you entered during the primary installation (default is asi).
 - b. When prompted for the Secondary Server URL, leave this blank, unless you plan on installing one of the Security Service Modules on the machine that hosts the backup Administration Application. In this case, specify the URL of the primary Administration Application.
 - c. When entering the Database Configuration, make sure to use the exact same configuration that you specified during the primary installation.
 - d. The database password used by the primary and backup Administration Applications are independent of each other. This allows you to use instance specific passwords to protect the various sensitive artifacts on the Administration Applications. For example, you may use different key passwords for the CA, Admin, SSM, and SCM identities that you entered in the primary installation. The same applies to the Identity, Peer, and Trust key store passwords.

Note: Do not install the database schema at the conclusion of the backup installation process. Do not run the `initialize_backup_trust.bat` (on Windows platforms) or `initialize_backup_trust.sh` (on Unix platforms) command after installing the backup.

4. Initialize the Backup Administration Application Trust Stores.

Before starting the backup Administration Application, you must synchronize the various trust stores used by the backup Administration Application with those of the primary. If this is not done, the backup Administration Application will not trust the Security Service Modules currently enrolled with the primary, and as a result, failover will not work.

To initialize the backup Administration Application trust stores, do the following:

- a. Copy the primary Administration Application and primary Service Control Manager ssl directories to the machine that hosts the backup Administration Application. One way to do this is to copy the primary Service Control Manager ssl directory onto a floppy disk or other removable media, and rename the ssl directory to scm-ssl. Likewise, copy the primary Administration Application ssl directory to the same removable media and rename the ssl directory to admin-ssl. Mount the removable media on the machine that hosts the backup Administration Application.

- b. Execute the `initialize_backup_trust.bat` (on Windows platforms) or `initialize_backup_trust.sh` (on Unix platforms) command from the `bin` directory of the backup Administration Application install. When prompted for the primary Service Control Manager SSL directory, enter the path to the `scm-ssl` directory created in the previous step. Likewise, when prompted for the primary Admin SSL directory, enter the path to the `admin-ssl` directory created in the previous step.
 - c. Delete the `ssl` directories on the removable media.
5. Start the backup Administration Application as you normally would a primary.
 6. Configure the Backup Server Trust Synchronization Mechanism.

Even though the backup Administration Application trust stores are synchronized with those of the primary in step 4, it is possible for them to become out of sync over time. This happens when a new Security Service Module or Service Control Manager is enrolled with the primary Administration Application. The trust stores of the primary Administration Application are updated with the new Security Service Module or Service Control Manager certificate during enrollment, but since an Security Service Module or Service Control Manager should only be enrolled with its primary, the backup Administration Application trust stores do not have the new certificate.

A similar trust situation occurs during Security Service Module or Service Control Manager unenrollment. To prevent the trust stores from becoming unsynchronized, the Administration Application has a trust synchronization mechanism that should be enabled on the backup Administration Application. The trust synchronization mechanism on the backup Administration Application periodically polls the primary for any updates to its trust store, and if a change has occurred, the mechanism updates the backup trust store with the contents of the primary. It is very important that you only enable the trust synchronization mechanism on the backup Administration Application.

See “Configuring the Administration Server for Failover” in the Administration Application Console Help for details on how to configure the backup server.

Installing

Post Installation Tasks

This section discusses the steps you need to take after installing the Administration Server.

Note: When installing, the administrator is usually logged in under a different account than the account on which the servers run when running as a service or daemon process. For this reason, it is important that the administrator ensure that the database client directories have appropriate permissions for the administration server user (`asiadmin` by default) to be able to access the database files. The inability of the administration server user to access the database files can result in services not being able to run or daemon processes or failing because they cannot access their database.

- [“Installing the Policy Database Schema” on page 5-2](#)
 - [“Installing the Policy Database Schema on Windows” on page 5-3](#)
 - [“Installing the Policy Database Schema on Sun Solaris” on page 5-4](#)
 - [“Installing the Policy Database Schema on Linux” on page 5-5](#)
- [“Starting and Stopping Processes” on page 5-6](#)
- [“Logging into the Administration Console” on page 5-6](#)
- [“Changing the System Password” on page 5-7](#)
- [“Fine Tuning your Application” on page 5-8](#)
- [“What’s Next?” on page 5-9](#)

Installing the Policy Database Schema

If you did not complete the final step in the installation program, installing the database schema, you must do that now. You only need to perform this step once. Before beginning this procedure, ensure that you have completed the following configuration and setup steps:

- Set the current `PATH` environment variables.

For Windows:

- Add product installation `lib` and `bin` directories (for example, `WLES_HOME\lib` and `WLES_HOME\bin`) to the `PATH` on Microsoft Windows.

For Sun Solaris and Linux:

- Add product installation `lib` and `bin` directories (for example, `WLES_HOME/lib` and `WLES_HOME/bin`) to `LD_LIBRARY_PATH` on Sun Solaris.

- Set the current `PATH` environment variables for your database.

- For Oracle:

Ensure that the Oracle client is set up and configured as described in [Chapter 3, “Setting Up and Administering the Database.”](#)

Ensure you can connect to the Oracle database server using command `sqlplus` (the Net Service Name, login ID and password).

For Windows, ensure that the `PATH` includes the `BIN` and `DLL` directory of the Oracle installation.

For Sun Solaris and Linux, ensure that the environmental variable `ORACLE_HOME` is set, `$ORACLE_HOME/bin` is in the `PATH`, and `$ORACLE_HOME/lib` is in the `LD_LIBRARY_PATH`.

- For Sybase:

Ensure that the Sybase 12.5 client is set up and configured as described in [Chapter 3, “Setting Up and Administering the Database.”](#)

In Windows, ensure that the `PATH` includes `%SYBASE%\OCS-12_5\bin` and `%SYBASE%\OCS-12_5\dll`. In Unix, ensure `PATH` includes `$(SYBASE)/OCS-12_5/bin`, and `LD_LIBRARY_PATH` includes `$(SYBASE)/OCS-12_5/lib`.

Ensure you can connect to the Sybase database server using command `isql` (the name of the database server, login ID and password).

Installing the Policy Database Schema on Windows

To install the policy database schema in a Microsoft Windows environment:

1. Change to the active directory in which to install the database schema, for example:

```
cd \bea\wles42-admin\bin
```

2. To install the database:

For an Oracle database, type:

```
install_schema_oracle.bat server dblogin dbpassword enterprise_domain
[policyowner]
```

For a Sybase database, type:

```
install_schema_sybase.bat server database dblogin dbpassword
enterprise_domain [policyowner]
```

Where:

server—The name of the Oracle net service name or Sybase server name.

database—The name of the Sybase database.

dblogin—The username to use to access the database; the username for the database administrator. Owner of the policy database (optional, defaults to the user login, usually the same as the *username*). The policy owner is a database user name or user ID that controls the database schema in the database instance.

dbpassword—Password to use to access the database; the password for the database administrator.

enterprise_domain—The name of the enterprise domain. The enterprise domain name is used to link all the components and is referred to as the Enterprise Domain Name when you installed the Administration Application.

[policyowner] —The Owner of the tables/schema in the policy database.

For more information on the database schema installation, examine the

`install_schema_oracle.log` or `install_schema_sybase.log` in the log directory.

Installing the Policy Database Schema on Sun Solaris

To install the policy database schema in a Sun Solaris platform:

1. Change to the active directory in which to install the database schema, for example:

```
cd /bea/wles42-admin/bin
```

2. Locate the script `install_schema_dbtype.sh`

Important: Make sure all scripts in this directory have execute permission.

3. To install the database:

For an Oracle database, type:

```
install_schema_oracle.sh server dblogin dbpassword enterprise_domain  
[policyowner]
```

For a Sybase database, type:

```
install_schema_sybase.sh server database dblogin dbpassword  
enterprise_domain [policyowner]
```

Where:

server—The name of the Oracle net service name or Sybase server name.

database—The name of the Sybase database.

dblogin—The username to use to access the database; the username for the database administrator. Owner of the policy database (optional, defaults to the user login, usually the same as the *username*). The policy owner is a database user name or user ID that controls the set of database schema in the database instance.

dbpassword—The password to use to access the database; the password for the database administrator.

enterprise_domain - Name of the enterprise domain. The enterprise domain name is used to link all the components and is referred to as the Enterprise Domain Name when you installed the Administration Application.

[policyowner] —The owner of the tables/schema in the policy database.

For more information on the database schema installation, examine the `install_schema_oracle.log` or `install_schema_sybase.log` in the log directory.

Installing the Policy Database Schema on Linux

To install the policy database schema in a Linux platform:

1. Change to the active directory in which to install the database schema, for example:

```
cd /bea/wles42-admin/bin
```

2. Locate the script `install_schema_dbtype.sh`

Important: Make sure all scripts in this directory have execute permission.

3. To install the database:

For an Oracle database, type:

```
install_schema_oracle.sh server dblogin dbpassword enterprise_domain
[policyowner]
```

For a Sybase database, type:

```
install_schema_sybase.sh server database dblogin dbpassword
enterprise_domain [policyowner]
```

Where:

server—The name of the Oracle net service name or Sybase server name.

database—The name of the Sybase database.

dblogin—The username to use to access the database; the username for the database administrator. Owner of the policy database (optional, defaults to the user login, usually the same as the *username*). The policy owner is a database user name or user ID that controls the set of database schema in the database instance.

dbpassword—The password to use to access the database; the password for the database administrator.

enterprise_domain—The name of the enterprise domain. The enterprise domain name is used to link all the components and is referred to as the Enterprise Domain Name when you installed the Administration Application.

[policyowner] —The owner of the tables/schema in the policy database.

For more information on the database schema installation, examine the

`install_schema_oracle.log` or `install_schema_sybase.log` in the log directory.

Starting and Stopping Processes

After you have installed the Administration Application, you must start the necessary processes by running the appropriate batch or shell scripts. On Windows, you can start these processes as services from the Programs menu or as commands from a console window.

For instructions on how to start and stop the required processes, see "[Starting and Stopping Processes](#)" in the *Administration Application Guide*.

Logging into the Administration Console

At this time, you can log into the Administration Console and check that all the components are working correctly. For descriptions of the process that is running, see "[Starting and Stopping Processes](#)" in the *Administration Application Guide*.

To start the Administration Console:

1. Open Internet Explorer.

To ensure that your transactions are securely encrypted, the Administration Console uses two-way Secure Socket Layers (SSL) to communicate with your Administration Server.

2. Open the URL for your Administration Console:

```
https://hostname:port/asi
```

Where:

`hostname` is the Domain Name Server (DNS) name or IP address of the Administration Server.

`port` is the port number through which the Administration Server is connected.

`asi` is the name of the Enterprise Domain (that you assigned during the installation procedure).

3. When the login page appears, enter the username and the password granted to one of the security roles with a login privilege and click Sign In. If you are using the default username and password, enter `system` (username) and `weblogic` (password). This is the default administrator configured on install and should only be used for the initial login.
4. Several security certificate verification dialog boxes appear. Check OK on each one. If you do not have the proper version of the JRE installed, then on the first attempt, the console prompts you to install it.

- Once you have started the console, you should set up additional administrative users or configure an Authentication provider to authenticate console users to an external authentication source such as LDAP or Microsoft Windows NT and update the administration policy accordingly, as described in [“Changing the System Password.”](#)

Note: The Administration Console allows administrators to edit configurations or perform other operations based on security roles granted by the administration policy. If your security roles do not permit editing of configuration data, for example, the data is displayed in the Administration Console but is not editable. If you try to perform an operation that is not permitted, the Administration Console displays an `Access Denied`.

Changing the System Password

During installation, a system username and password are defined for use when you first start the console. To change the system password, you must change it in the Administration Console and `boot.properties` file and run the `asipasswd` Utility (`asipassword.bat`) to update the `password.xml` and `password.key` files.

Note: BEA recommends changing the `system` password that was set during installation. Each Administration Application deployment must have a unique password. For additional information on how to begin using the Administration Console, see the online help.

To change the `system` password, perform the following steps:

- Start the Administration Console and open the Identity folder.
The Identity page displays the name of each directory available.
- Select the `WLES` directory, click `Users`, and then select `System`.
The Users page displays one user named `System`.
- Click `Edit` and click `Set Password`.
- In the `Old Password` text box type the current password. (The default password set during installation is `weblogic`.)
- In the `Password` text box, type the new password for the `System` user, confirm the password in the `Confirm` text box, and click **OK**.
- Open the `boot.properties` file located at `BEA_HOME\wles42-admin\asiDomain`, delete the encrypted text in the `username` and `password` fields, enter `system` in the `username` field and the new password in the `password` field, and save the file.

7. To update the system password in the `password.xml` file, perform the following steps to run the `asipasswd` Utility:
 - a. Open a command window, go to `BEA_HOME\wles42-admin\bin`, and enter the following command:

```
asipassword.bat system ..\ssl\password.xml ..\ssl\password.key
```
 - b. When prompted, enter and confirm the new password.
8. Restart the Administration Application server and login into the Administration Console using the new password.

Fine Tuning your Application

WebLogic Enterprise Security provides certain security properties that control the behavior of the Web Service client bindings and socket pooling routines. These properties effect how WebLogic Enterprise Security performs under load and allow you to fine tune the machine of which the application is running.

To improve performance, you can configure the Web Service client HTTP/HTTPS bindings and socket pooling routines to control how the socket pooling behaves. These properties can either be set in a file called `security.properties` located on the local machine, in the working directory, or can be passed as Java system properties using the command line (-D) argument.

- `wles.webservice.pool.MinimumPoolSize`

Sets the minimum size of the socket pool. The socket pooling routines create new sockets if the size of the pool ever drops below this number. The default value is 0.

- `wles.webservice.pool.InitialPoolSize`

Sets the initial size of the socket pool. At creation time, the pool populates the socket pool with this number of connections. The default value is 0.

- `wles.webservice.pool.MaximumPoolSize`

Sets the maximum size of the socket pool. The socket pooling routines ensure that the pool never grows larger than this number. The default value is the number of active threads in the JVM at pool creation time.

- `wles.webservice.pool.PoolDeltaSize`

Sets the number of sockets that can be added to or removed from the pool at a given time. The default value is 1.

- `wles.webservice.pool.SocketInactivityTimeout`

Sets the amount of time (in seconds) before an inactive socket is eligible for eviction from the socket pool.

- `wles.webservice.http.RequestTimeout`

Sets the amount of time (in milliseconds) before a blocking that a Web Service read operation generates a timeout. If the value of this property is 0, read operations block indefinitely. The default value is 0.

- `wles.webservice.http.SocketPoolingEnabled`

Disables or enables socket pooling. If disabled, a Web Service client creates a new socket for each request. The default value is true.

Warning: BEA strongly discourages the disabling of socket pooling.

What's Next?

Now that you have successfully installed the Administration Application, you are ready to install your Security Service Modules and deploy your security configurations and policies.

For instructions on how to install Security Service Modules, see the following documents:

- [WebLogic 8.1 Security Service Module \(SSM\) Installation](#)
- [Java SSM Installation](#)
- [Web Server SSM Installation](#)

Note: In a production environment, BEA recommends that you install your Security Service Modules on machines other than the machine on which the Administration Server is installed.

For instructions on how to design and write security policy to protect resources, see the [BEA WebLogic Enterprise Security Policy Managers Guide](#). This document defines the policy model used by BEA WebLogic Enterprise Security, and describes how to import and export policy data.

For instruction on how to setup a metadirectory to extract user data from your user repository, see [“Configuring Metadirectories” on page 6-1](#).

Post Installation Tasks

Configuring Metadirectories

This section describes how to configure a metadirectory to extract user data from your user repository (for example, an LDAP server, an Active Directory, a database server, or NT Domain directory) and import that data into the policy database. As a result, the user, group and attribute data (referred to simply as attributes) are available and synchronized, and can be used to enforce dynamic security policies in your applications through the ASI Authorization and ASI Role Mapping services.

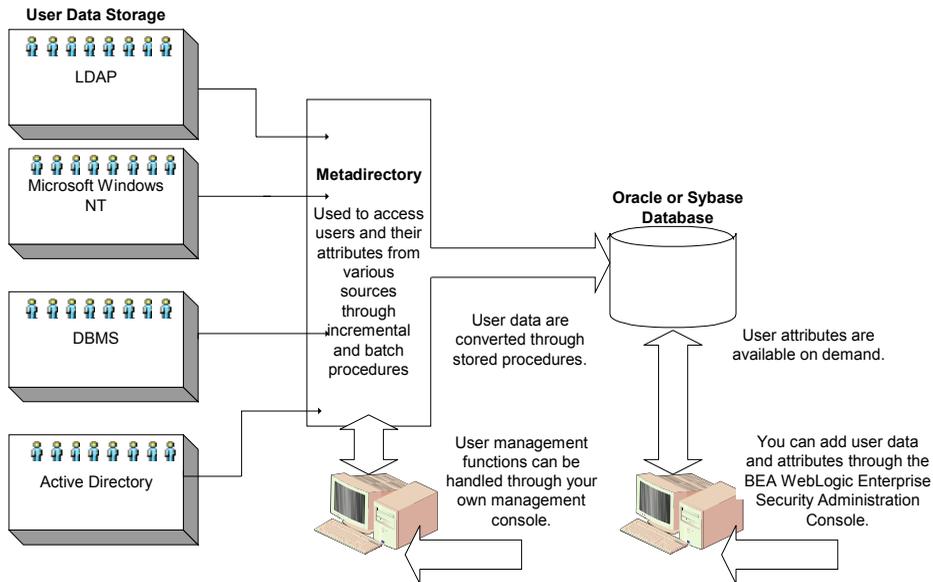
This section covers the following topics:

- [“Why Use Metadirectories?” on page 6-2](#)
- [“Metadirectory Configuration Overview” on page 6-3](#)
- [“Preparing to Configure a Metadirectory” on page 6-4](#)
- [“Configuring Metadirectory Tables and Database Triggers” on page 6-5](#)
- [“Configuring Metadirectory Schemas” on page 6-13](#)
- [“Configuring Metadirectory Synchronization” on page 6-28](#)
- [“Verifying that Metadirectory Synchronization Works” on page 6-34](#)

Why Use Metadirectories?

BEA WebLogic Enterprise Security requires that all policy data be stored in either an Oracle or Sybase policy database. The goal of a metadirectory is to provide your organization with a unified view of all identity information. A metadirectory solves important business issues that result from having information stored in multiple, disparate data repositories throughout an organization. Thus, through the use of a metadirectory, the maintenance cost of sharing information is reduced and the accuracy and the overall security of an application is improved (see [Figure 6-1](#)).

Figure 6-1 Metadirectory Architecture



In BEA WebLogic Enterprise Security, the term directory applies to any collection of user data, stored in a database, LDAP directory server, or other type of repository. These directories form the core of any identity management solution because every user repository has its own approach to the storage of information.

An identity directory refers to any user repository configured for use with BEA WebLogic Enterprise Security. In the Administration Console, the identity directory defines a logical collection of users, groups and attributes that can be used to design your authorization and role mapping policy, and store information about who is authorized through your policies. An identity

directory typically represents groups of users of a particular application or resource, users in a specific location, or users imported from an external user repository.

A metadirectory can be used to store attributes replicated and synchronized from your user repository into the policy database. Following replication, the user data are available as attributes through the Administration Console identity directories for use in your authorization policies and policy rules. Any changes that you make to the replicated user attributes using the Administration Console are not propagated back to your user repository.

Metadirectory Configuration Overview

To configure metadirectories, you use several different components (see [Figure 6-2](#)). In [Figure 6-2](#), each task is represented by a number that is positioned next to the component that you use to perform the task.

Figure 6-2 Metadirectory Configuration Components

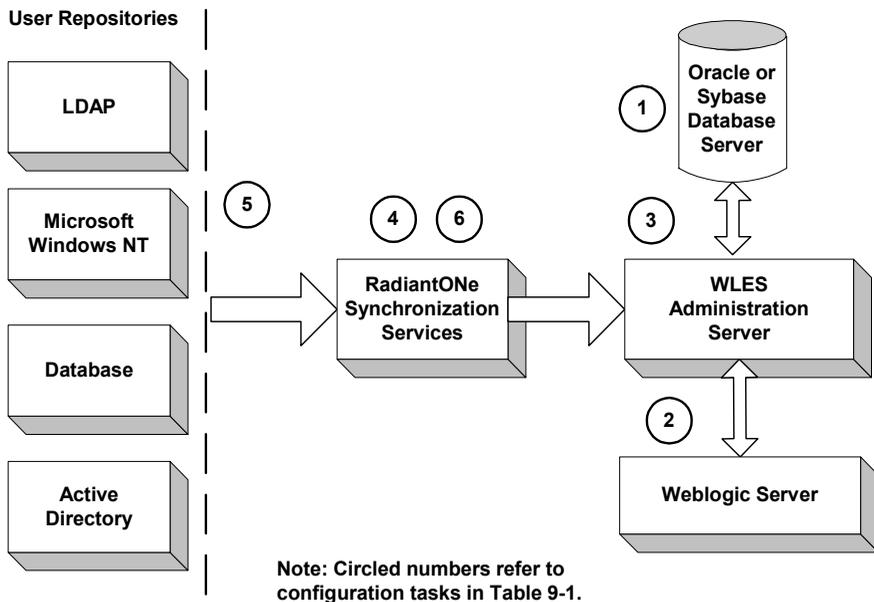


Table 6-1 summarizes the tasks and links each task to a circled number in Figure 6-2.

Table 6-1 Metadirectory Configuration Tasks

Use this component:	To perform these tasks:
Database Server	1. Create the destination tables.
WebLogic Server Administration Console	2. Configure a JDBC connection pool and the Java Message Service.
WebLogic Enterprise Security Administration Console	3. Configure database triggers.
RadiantOne Synchronization Services	4. Configure metadirectory schemas and the synchronization hub.
User Repository Server	5. Configure the directory connector. Note: This task is required only if you are using Sun ONE Active Directory as your user repository server.
RadiantOne Synchronization Services	6. Configure the connectors and starting the synchronization hub and the connectors.

For detailed instructions on performing each of the tasks listed in Table 6-1, see the following sections:

- [“Preparing to Configure a Metadirectory” on page 6-4](#)
- [“Configuring Metadirectory Tables and Database Triggers” on page 6-5](#)
- [“Configuring Metadirectory Schemas” on page 6-13](#)
- [“Configuring Metadirectory Synchronization” on page 6-28](#)

Preparing to Configure a Metadirectory

Before you begin, you must install and start the RadiantOne Synchronization Services. RadiantOne Synchronization Services use the JDBC and Java Message Service (JMS) features of the WebLogic Server that hosts the WebLogic Enterprise Security Administration Application to update the metadirectory. The RadiantOne Synchronization Services tool provides connectors for the master identity repositories that send XML formatted messages whenever information in a user repository is updated. Through the use of the JMS, this tool ensures that all updates are

delivered and processed. For installation instructions, see [“Installing the RadiantOne Synchronization Services” on page 6-5](#).

Installing the RadiantOne Synchronization Services

The RadiantOne Synchronization Services software is available as separate installation CD-ROMs (see install kit disks 3 and 4). After you install the product, you can access the applications from the Start>Programs>RadiantOne menu.

To install the RadiantOne Synchronization software:

- On a Microsoft Windows platform, run: wles421metadir_win32.exe.
- On a Sun Solaris platform, run: wles421metadir_solaris32.bin.
- On a Linux platform, run: wles421metadir_solaris32.bin.

Note: A second installation package is available for installing the RadiantOne Connectors. These are included in the RadiantOne Synchronization Services install package and do not need to be installed separately. However, you can install the connector package separately on another machine. For instructions for installing the connectors separately, refer to the RadiantOne Synchronization Services installation documentation. When you install the RadiantOne software, the online documentation is installed in the RadiantOne directory

Configuring Metadirectory Tables and Database Triggers

This section covers the tasks that you must perform to create destination tables in the WebLogic Enterprise Security policy database and install triggers on those tables.

To configure metadirectory tables and database triggers, perform the following tasks:

- [“Creating Metadirectory Destination Tables” on page 6-5](#)
- [“Configuring a JDBC Connection Pool and JMS” on page 6-9](#)
- [“Configuring Metadirectory Database Triggers” on page 6-11](#)

Creating Metadirectory Destination Tables

There are two tables that you have to create in the policy database for the synchronization of users and groups to work properly: `ASI_USERS` and `ASI_GROUPS`. After you create these tables, you configure them through the WebLogic Enterprise Security Administration Console. You must create these tables before you perform the remaining tasks in this section.

The following sections provide guidelines and restrictions for the tables and detailed instructions on how to create them:

- [“Metadirectory Destination Table Guidelines and Restrictions” on page 6-6](#)
- [“Creating Metadirectory Destination Tables Using Oracle or Sybase” on page 6-8](#)

Metadirectory Destination Table Guidelines and Restrictions

Two tables are required in the policy database for the synchronization of user repositories: `ASI_USERS` and `ASI_GROUPS`. The following sections provide guidelines and restrictions for these tables:

- [“User Synchronization Table Guidelines” on page 6-6](#)
- [“User Synchronization Table Guidelines” on page 6-6](#)
- [“User and Group Attributes Character Set Restrictions” on page 6-7](#)

User Synchronization Table Guidelines

The user synchronization table is used by the partner tool to stage user and user attribute information for import into the policy database. The name of the table used for user synchronization is configurable.

While the names of the columns in the table are configurable, the following restrictions apply:

- One column in the table must serve as the unique identifier for the user. The UID may contain any character, but the `'` character must be escaped. For example, `“John\Doe”` must be entered as `“John/Doe”`.
- The Primary Key for the table should be the column used as the UID. For performance and data consistency, the user synchronization table should include the primary key in its definition.

The user synchronization table accommodates source repositories that store group memberships as user attributes. Managing group memberships as user attributes does not impact managing group memberships explicitly through the group synchronization table—both ways can be used.

You must adhere to the following restrictions and requirements when setting up group memberships through the user synchronization table:

- Only one column may be used for storing the group memberships.
- The group column needs to be a character string (typically in Oracle: `varchar2`).

- Membership in multiple groups is possible and is stored as a delimited text string. The choice of delimiter is configurable but should be sufficiently uncommon so that parsing of the group list may be done correctly.
- If the group name contains the '/' character, it should be escaped.

Any number of columns in the user synchronization table may be used for passing attributes into the WebLogic Enterprise Security Administration Server. The columns used for all attributes in the User Synchronization table must be of variable length character (for example, in Oracle: varchar2). For purposes of importing from the user synchronization table, you may map attributes to any of the following WebLogic Enterprise Security policy data types: `string`, `integer`, `boolean`, `date`, `time`, `dayofweek_type`, `month_type`, and `object_type`. Attributes are also defined as either `list` or `single`. Multiple attribute values of type `list` are stored as a delimited text string. The delimiter used for attributes of type `list` must be the same as the delimiter used for groups.

Group Synchronization Table Guidelines

In addition to the user-attribute-based group membership discussed above, group memberships may also be defined by using the group synchronization table. Unlike the `user_synchronization` table, the schema for the group synchronization table is fixed, that is, it must adhere to the structure shown in [Table 6-2](#).

Table 6-2 Group Synchronization Table

Column Name	Type	Description
CN	varchar2	The name of the group
UNIQUEMEMBER	varchar2	The name of the user belonging to the group.

You must adhere the following restrictions and requirements when setting up the group synchronization table:

- The Primary Key for the table should consist of both columns.
- A forward slash (/) in the value for either of the columns must be escaped using a back slash (\).

User and Group Attributes Character Set Restrictions

The following requirements and restrictions apply to user and group attributes

- The name of the attribute cannot be longer than 1000 character (580 characters for some Sybase 12.5 configurations, depending on the page size)
- Each value of a user attribute cannot be longer than 1000 characters. (580 characters for some Sybase 12.5 configuration, depending on the page size)
- The length of the value of all user attributes combined cannot be longer than the lesser of 16,000 characters or the `varchar2` column-size limit for the database.
- Attribute names in WebLogic Enterprise Security may only consist of alphanumeric characters (a-z, A-Z, 0-9) and the underscore (`_`) character.
- The column name of the user synchronization table is limited by any database character set limitations.
- Attribute names must start with an alphabetic character or an underscore.
- Any printable characters are allowed except double quote (`"`) and back slash (`\`).

Creating Metadirectory Destination Tables Using Oracle or Sybase

To create the `ASI_USERS` and `ASI_GROUPS` destination tables using an Oracle or a Sybase database server,

1. To log into the policy database, open a command window and type:

```
sqlplus username/password@asi
```

where: `username` and `password` are the username and password you defined when you created the database user account and `asi` is the database instance name.

2. To create the `ASI_USERS` destination table, enter the following SQL command:

```
SQL> CREATE TABLE ASI_USERS(DisplayName VARCHAR(255) NULL,  
COMMONDOMAIN VARCHAR(255) NOT NULL, PRIMARY KEY (COMMONDOMAIN))
```

3. To create the `ASI_GROUPS` destination table, enter the following SQL command:

```
SQL> CREATE TABLE ASI_GROUPS(CN VARCHAR(255) NOT NULL,  
UNIQUEMEMBER VARCHAR(255) NOT NULL, PRIMARY KEY  
(CN, UNIQUEMEMBER))
```

Note: In addition to `DisplayName`, you can add more columns to the destination tables to be used as user attributes, such as street address, zip code, email, phone, and so on.

Configuring a JDBC Connection Pool and JMS

To connect to the RadiantOne Synchronization Services to the WebLogic Enterprise Security asiDomain, you must use the WebLogic Server Administration Console to configure a JDBC connection Pool and the Java Message Service (JMS).

To configure the JDBC connection pool and JMS, perform the following steps:

1. To start the WebLogic Server Administration Console, open a browser and go to `https://hostname:7010/console`,
where:
hostname in the name of the machine that is hosting the WebLogic Enterprise Security Administration Application
7010 is the port on which the Administration Console is running
2. In the left pane of the Administration Console, open the Services and JDBC folders and click Connection Pools. The asiDomain> JDBC Connection Pools page is displayed in the right pane.
3. Click Configure a new JDBC Connection Pool. The Configure a JDBC Connection Pool: Choose database page is displayed.
4. Select the Database Type and the Database Driver as specified in [Table 6-3](#) and click Continue. The Configure a JDBC Connection Pool: Define connection properties page is displayed.

Table 6-3 Database Type and Database Driver Parameter Settings

JDBC Connection Pool Parameter	Setting
Database Type	Oracle or Sybase
Database Driver	For Oracle 8i, select Oracle's Driver (Thin) Versions: 8.1.7 For Oracle 9i, select Oracle's Driver (Thin) Versions: 9.0.1,9.2.0,10 For Sybase, select BEA's Sybase Driver (Type 4) Versions: 11.X,12.X

5. Refer to [Table 6-4](#), enter the appropriate values in the Configure a JDBC Connection Pool: Define connection properties page, and click Continue. The Configure a JDBC Connection Pool: Test database connection page is displayed.

Table 6-4 JDBC Connection Pool Configuration Parameters

Parameter	Description
Name	The JDBC connection pool name that you specify, for example, ConsolePool.
Database name	The name assigned to the instance of the database when it was created, for example, ASI5
Hostname	The name of the machine on which the database server is installed, for example, ASI_host
Port	The port used for the connection to the database server (default: 1521).
Database User Name	The username of the database account, for example, wles
Password/Confirm Password	The password assigned with the database account was create for the user (any alphanumeric string).

6. Click Test Driver Configuration. A “Connection successful” message and the Configure a JDBC Connection Pool: Create and deploy page is displayed.
7. Click Create and Deploy. The connection pool is deployed.
8. To configure a JMS template, perform these steps:
 - a. In the left pane, open the Services and JMS folders and click Templates. The asiDomain> JMS Templates configuration page is displayed in the right pane.
 - b. Click Configure a new JMS Template, name the template RLI_JMS, and then click Create.
9. To configure a JMS JDBC store, perform these steps:
 - a. In the left pane, click Stores.
 - b. Click Configure a new JMS JDBC Store.
 - c. Name the store RLI_JDBC_STORE.
 - d. Set the Connection Pool to the name of the connection pool created previously (ConsolePool) and select Create.
10. To configure a JMS server, perform these steps:
 - a. In the left pane, click Servers.

- b. Click Configure a new JMS Server.
 - c. Name the server `RLI_JMS_SERVER`.
 - d. Set Persistent Store to the JDBC store that was created previously (`RLI_JDBC_STORE`).
 - e. Set Temporary Template to the template that was created previously (`RLI_JMS`), and click Create.
 - f. Click the Target an Deploy tab and set Target to `asiAdminServer` and click Apply.
11. To configure a JMS Connection factory, perform these steps:
- a. In the left pane, click Connection Factories.
 - b. Click Configure a new JMS Connection Factory,
 - c. Set Name to `RLI_JMS_CONNECTION`,
 - d. Set JNDI Name to `weblogic.asiAdminServer.jms.TopicConnectionFactory`, and click Create.
 - e. Click the Target an Deploy tab and set Target to `asiAdminServer` and click Apply.
12. To restart the WebLogic Server so that the change takes effect, close the WebLogic Server command window or, if WebLogic Server is setup to run as a Windows service, restart the service.
13. This completes the configuration of the JDBC connection pool and JMS.

Configuring Metadirectory Database Triggers

You must configure database triggers for the user and group synchronization tables in the policy database. A database trigger provides a necessary link between the metadirectory database and the policy database. A trigger enables the user attributes to be received by the Administration Server and put into the identity directory (that you define) whenever a change occurs in the underlying metadirectory database.

Note: Any modifications that you make to the existing data records in the synchronization tables must be made with an `UPDATE` command, not through a series of `DELETE` and `INSERT` commands. Use `INSERT` only for new records and `DELETE` only for removing records. Also, do not use the “truncate table” command to clean either the user or group synchronization tables because that command does not activate the triggers.

To configure metadirectory database triggers, perform the following steps:

Configuring Metadirectories

1. To start the WebLogic Enterprise Security Administration Console, open a browser and go to `https://hostname:7010/asi`,
where:
hostname is the name of the machine that is hosting the WebLogic Enterprise Security Administration Application
7010 is port on which the Administration Console is running
asi is the domain name
2. In the left pane, open the Identity folder and click Metadirectory Configuration. The Metadirectory Configuration page is displayed in the right pane.
3. In the Metadirectory Configuration page, select the database type (either Oracle or Sybase), enter the name of the JDBC connection pool (for example: ConsolePool) and the name of the synchronization tables (ASI_USERS and ASI_GROUPS), and click Connect. A “Successful Connection” message is displayed along with additional fields that require input.
4. Refer to [Figure 6-3](#), and fill in the additional fields. Set `user id` to the WebLogic Enterprise Security schema owner, which is the same as the database account username. Set the `identity directory name` field to any directory name that is unique in the `asiDomain`. This identity directory is the directory in the policy database into which users are populated. Set the `COMMONDOMAIN` and `DISPLAYNAME` parameters as shown in [Figure 6-3](#).
5. Click Install Trigger. A “Trigger successfully installed” message is displayed.
6. This completes the configuration of the metadirectory database triggers.

Figure 6-3 Metadirectory Triggers Configuration

Connecting to Database Connection Successful

Enter the user id of the owner of the WebLogic Enterprise Security database schema.

Enter the name of the Identity Directory that will contain the synchronized users in the metadirectory.

Enter the delimiter used to separate multi-value attributes

Name	Configuration Type	WLES Attribute Type	Single or List
COMMONDOMAIN	<input type="text" value="uid"/>	<input type="text" value="string"/>	<input type="text" value="Single"/>
DISPLAYNAME	<input type="text" value="attribute"/>	<input type="text" value="string"/>	<input type="text" value="List"/>

Configuring Metadirectory Schemas

BEA WebLogic Enterprise Security uses a comprehensive schema for tracking and updating all policy data. You use RadiantOne Synchronization Services to configure the schemas required to upload user and groups information from the user repository to the policy database.

To configure the required metadirectory schemas, perform following tasks:

- [“Extracting the Source Schemas” on page 6-14](#)
- [“Loading the Source Schemas” on page 6-17](#)
- [“Extracting the Destination Schemas” on page 6-19](#)
- [“Loading the Destination Schemas” on page 6-22](#)
- [“Configuring the Source-to-Destination Topology” on page 6-23](#)
- [“Configuring the Topology Transformations” on page 6-25](#)
- [“Uploading User and Group Data” on page 6-26](#)

Extracting the Source Schemas

This section describes how to extract the source schemas for the user repository.

To extract the source schemas from the user repository, perform the following steps:

1. Copy the `WL_Home\server\lib\weblogic.jar` file to the `RadiantOne\r1syncsvcs\bea_lib` directory.
2. To start the RadiantOne Synchronization Services tool:
 - On Windows: click `Start>Programs>RadiantOne>RadiantOne Synchronization Services>Synchronization Services Administrator`.
 - On Sun Solaris: From the `RadiantOne/r1syncsvcs/bin` directory, run: `runSSC.sh`.
3. To start the Schema Extraction Wizard is displayed, select `New` from the `Datastore` drop-down menu.
4. Select `LDAP Schema Extraction` radio button and click `Next`. The `LDAP Schema Extractor` page is displayed.
5. Refer to [Table 6-4](#) and [Table 6-5](#) and enter the directory server information. To determined the complete directory server name and port number, refer to the directory server console (see [Figure 6-5](#)) and check the values. For example, in [Figure 6-5](#), the complete name is `asi_host.amer.bea.com` and the port number is `56763`.

Figure 6-4 LDAP Schema Extractor Page

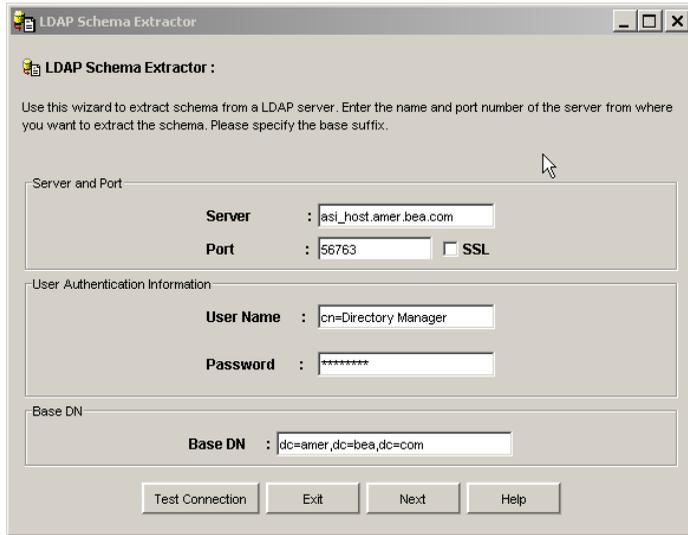


Table 6-5 LDAP Schema Extractor Parameters

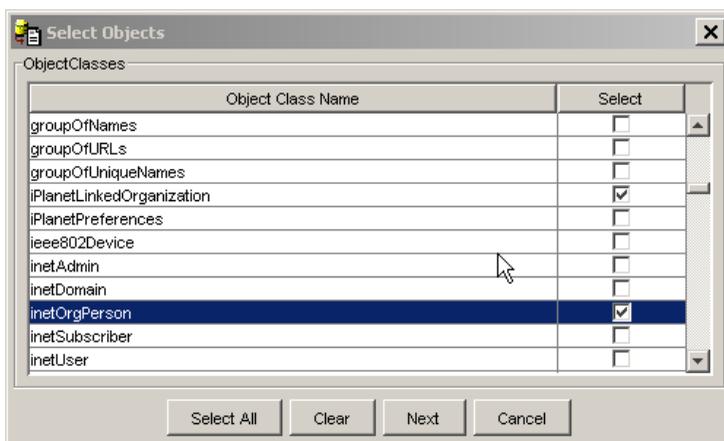
Parameter	Description
Server	The name of the LDAP Directory server, for example, <code>asi_host.amer.bea.com</code>
Port	The port number of the LDAP Directory server, for example, <code>56763</code>
Username	The username you enter to access the LDAP Directory server (default: Directory Manager).
Password	The password you enter to access the LDAP Directory server.
Base DN	The base domain name, for example, <code>dc=amer, dc=bea, dc=com</code> .

Figure 6-5 Sun ONE Server Console



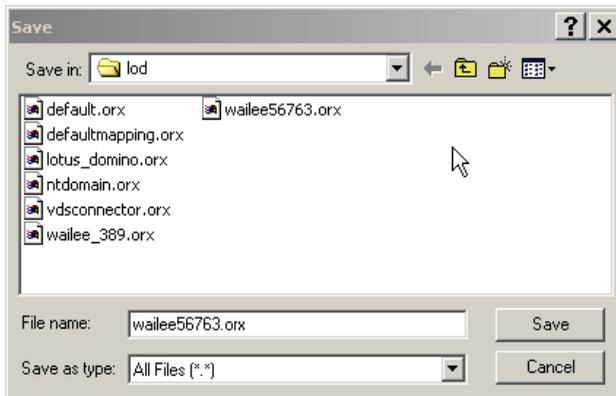
6. Click Test Connection. A “Connection Successful” message dialog box is displayed.
7. To close the message dialog box, click Ok and then click Next. The Select Objects page is displayed (see Figure 6-6).

Figure 6-6 RadiantOne Select Objects Page



8. Select the `groupOfUniqueNames` and `inetOrgPerson` object classes and click Next. The Save windows is displayed (see [Figure 6-7](#)).

Figure 6-7 RadiantOne Select Object Save Window



9. In the Save window, edit the Filename field to remove all but the directory server name, the port number, and the `.orx` filename extension as shown in [Figure 6-7](#), and click Save. A “Schema Extraction Completed” message dialog box is displayed.
10. . Click Ok and then click Exit.
11. This completes the extraction of the source schemas.

Loading the Source Schemas

This section describes how to load the source schemas for the user repository.

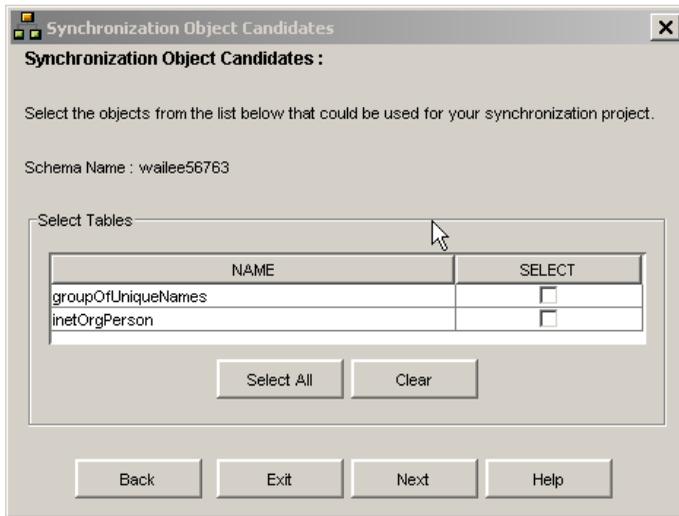
To load the source schemas from the user repository, perform the following steps:

1. On the RadiantOne Synchronization Services Administration console, from the Datastore drop-down menu, select Add. The Synchronization Object Candidates page is displayed (see [Figure 6-8](#)).

Figure 6-8 RadiantOne Synchronization Object Candidates Page



2. Click Directory. The Open window is displayed.
3. Select schema extraction file that you created in “[Extracting the Source Schemas](#)” on [page 6-14](#) (for example: `asi56763.orx`) and click Open. The Schema File field is populated with the filename, including the path.
4. Click Next. The Synchronization Objects Candidates page is displayed (see [Figure 6-9](#))

Figure 6-9 Synchronization Select Objects Page

5. Click Select All and click Next. A “Successfully generated the datastore” message is displayed.
6. To exit, click Finish.
7. This completes the loading of the source schemas.

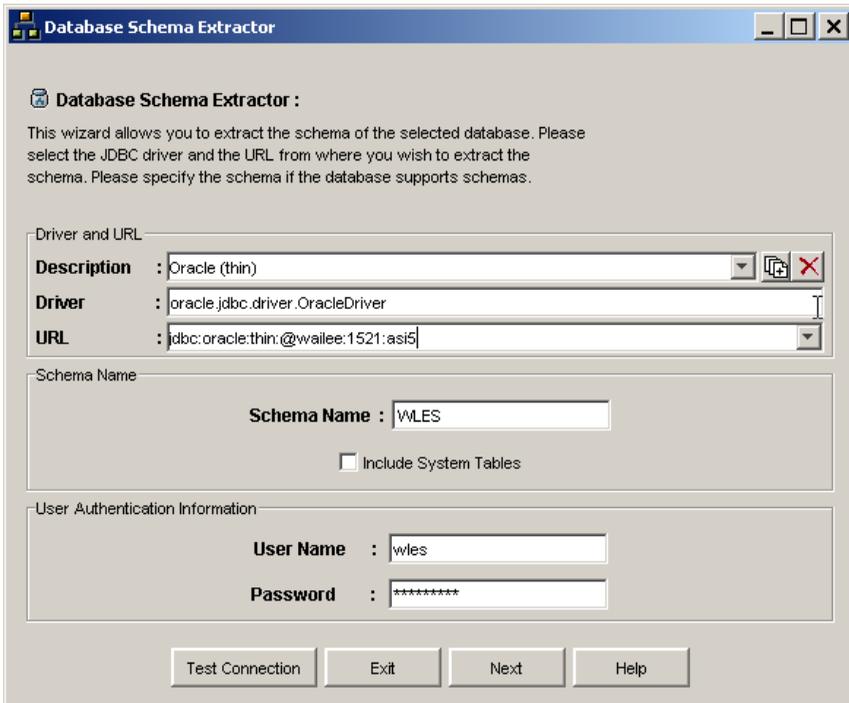
Extracting the Destination Schemas

This section describes how to extract the policy database destination schemas for the database server.

To extract the destination schemas from the database server, perform the following steps:

1. On the RadiantOne Synchronization Services Administration console, to start the Schema Extraction Wizard is displayed, select New from the Datastore drop-down menu.
2. Select Database Schema Extraction radio button and click Next. The Database Schema Extractor page is displayed (see [Figure 6-10](#)).

Figure 6-10 Database Schema Extractor Page



3. Refer to [Table 6-6](#) and set the database server parameters.

Table 6-6 Database Server Parameters

Parameters	Description
Description	The type of database used for the database server.
Driver	The database JDBC driver. Note: Better performance may be achieved by configuring a Type II JDBC driver. For Oracle (OCI), this is the same driver but uses a different URL syntax. Please refer to your Oracle documentation for the correct syntax and configuration.

Table 6-6 Database Server Parameters (Continued)

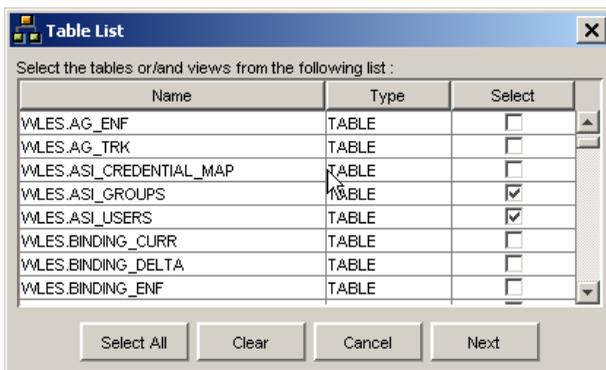
Parameters	Description
URL	The URL of the database server. You are only required to supply the name of the database host machine and database name, for example, jdbc:oracle:thin@asi_host:1521:asi5.
Schema Name	The schema name. This name must be unique in the database server, for example, WLBS. Note: If you are using an Oracle database server, you must type the schema name in uppercase.
Include System Tables	Determines whether system files are included. Check this box to on.
User Name	The username you enter to access the database server.
Password	The password you enter to access the database server

- To verify that the schema extractor can connect to the database server, click Test Connection. A “Connection succeeded” dialog box is displayed.

Note: If the connection fails, make sure that all the database server parameters are set correctly.

- Click Ok and click Next. The Table List dialog box is displayed (see [Figure 6-11](#)).

Figure 6-11 Table List Dialog Box



6. Select the `WLES.ASI_GROUPS` and `WLES.ASI_USERS` tables and click Next. The Save window is displayed.
7. Accept the default filename (the default filename should match the database name) and click Open. A “Schema Extraction Completed” message dialog box is displayed.
Note: Be sure to save the filename in lowercase. Also, the filename cannot contain any periods (for example, this filename is correct: `asi5.orx`).
8. Click Ok and then click Exit. By default this schema is saved to `../RLI_HOME/data/org`.
9. This completes the extraction of the destination schemas.

Loading the Destination Schemas

This section describes how to load the policy database destination schemas for the database server.

To load the destination schemas from the database server, perform the following steps:

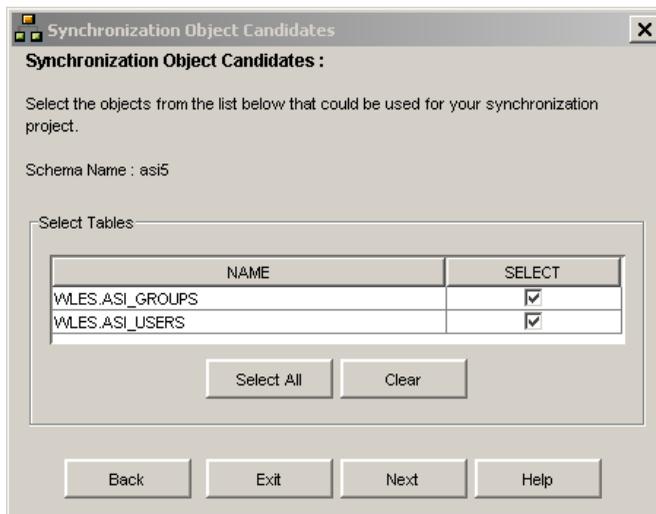
1. On the RadiantOne Synchronization Services Administration console, from the Datastore drop-down menu, select Add. The Synchronization Object Candidates page is displayed (see [Figure 6-12](#)).

Figure 6-12 RadiantOne Synchronization Object Candidates Page



2. Click Database. The Open window is displayed.
3. Select schema extraction file that you created in “[Extracting the Destination Schemas](#)” on [page 6-19](#) (for example: `asi5.orx`) and click Open. The Schema File field is populated with the filename, including the path.
4. Click Next. The Synchronization Objects Candidates page is displayed (see [Figure 6-13](#)).

Figure 6-13 Synchronization Select Objects Page (Database Server Objects)



5. Click Select All and click Next. A “Successfully generated the datastore” message is displayed.
6. To exit, click Finish.
7. This completes the loading of the database server schemas for the policy database.

Configuring the Source-to-Destination Topology

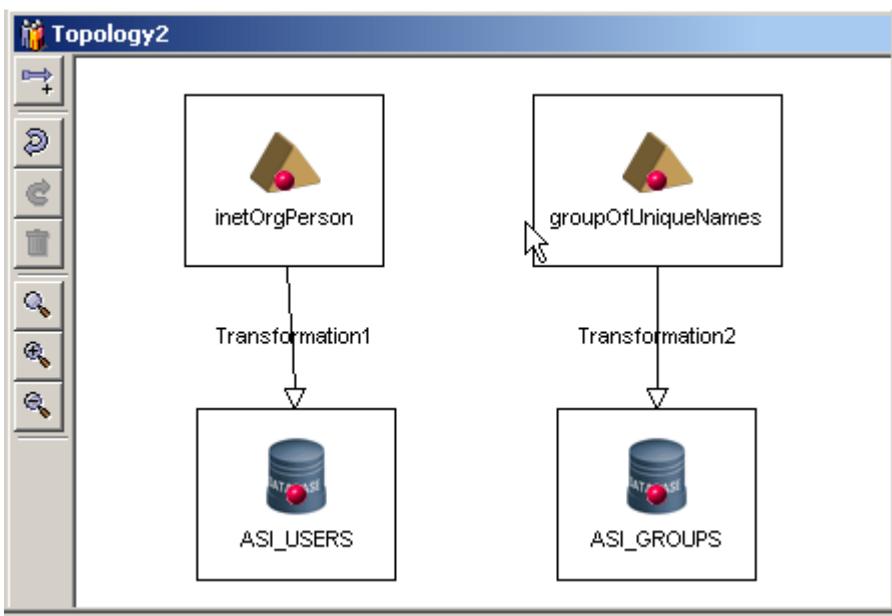
This section describes how to configure the RadiantOne topology that serves to link the source user repository to the policy database. The topology shows all of the objects involved in the synchronization process and the data flow. You use the topology to define and connect all of the data objects that are involved in a particular synchronization process.

To configure the topology, perform the following steps:

Configuring Metadirectories

1. On the RadiantOne Synchronization Services Administration console, from the RadiantOne Synchronization Services Topology drop-down menu, select New. The Topology window is displayed in the right pane.
2. Expand the nodes in the left pane, click the `groupOfUniqueNames` node and drag and drop it into the right pane.
3. Repeat step 3 for the `inetOrgPerson`, `ASI_GROUPS`, and `ASI_USERS` nodes.
4. Click the red dot for each of the publishing objects (`groupOfUniqueNames` and `inetOrgPerson`) and drag it to the red dot of the subscribing object. The tool draws lines connecting to the objects and labels them Transformation1 and Transformation2 (see [Figure 6-14](#)).

Figure 6-14 Topology Layout



Configuring the Topology Transformations

This section describes how to configure the RadiantOne topology transformation scripts. These scripts determine how the source data in the user repository is transformed before it is written to the policy database.

Note: The RadiantOne Synchronization Services tool is capable of very sophisticated transformations, including creating a unified identity from multiple sources. The transformation scripting language is Java. If you want to explore more sophisticated transformations or merging of identity data, refer to the RadiantOne documentation available in the RadiantOne installation directory.

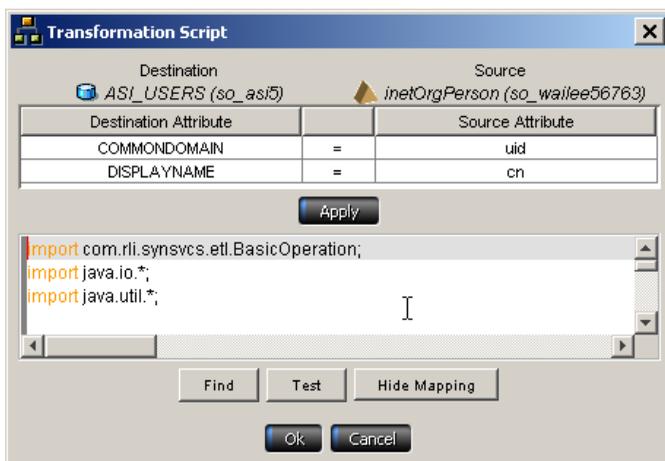
The BEA WebLogic Enterprise Security product ships with sample user and group transformation scripts. They are located at

`BEA_Home\wles42-admin\examples\r1syncservice`. The file names are `asi_users.djava` and `asi_groups.djava`. The samples are dynamic java scripts that are compiled and run by the RadiantOne Synchronization Services as part of its transformation runtime. The following procedure uses the `asi_groups.djava` sample script.

To configure the topology transformation scripts, perform the following steps:

1. On the RadiantOne Synchronization Services Administration console, right click Transformation1 and select Edit Script. The Transformation Script window is displayed (see [Figure 6-15](#)).

Figure 6-15 Topology Script for Transformation1



2. Set the COMMONDOMAIN and DISPLAYNAME attributes to uid and cn respectively as shown in [Figure 6-15](#), click Apply, and click Ok.
3. Right click Transformation2 and select Edit Script. The Transformation Script window is displayed.
4. Position your cursor in the bottom region of the window, right click and select load. The Open window is displayed.
5. Locate the asi_groups.djava file in the BEA_Home\wles42-admin\examples\r1syncservice directory, select it and click Open, click Apply, and click Ok.
6. Click the Topology drop-down menu and click Save to save the topology.
7. This completes the configuration of the transformation topology.

UpLoading User and Group Data

This section describes how to use the transformation topology to upload the user and group data from the source user repository to the policy database.

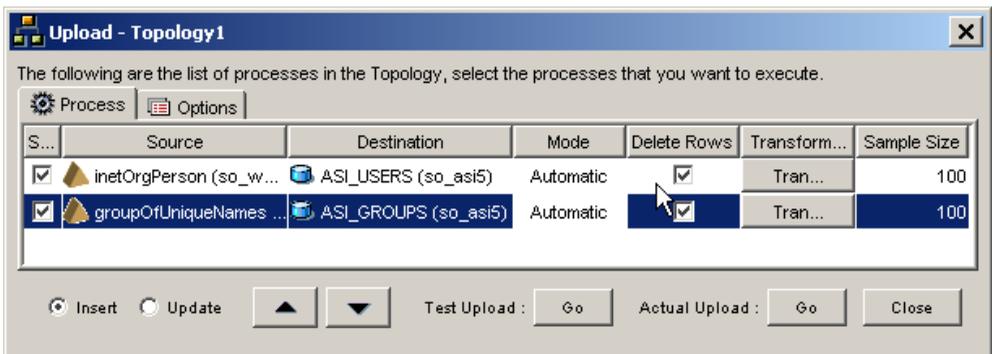
Note: This task serves to test all that all the configuration tasks you have performed up to this point have been performed correctly and that you can proceed to the next section, [“Configuring Metadirectory Synchronization” on page 6-28](#), and perform the

synchronization tasks. If the upload fails, check the previous tasks to ensure that they were performed correctly. Also, verify that you used the correct passwords in each of the previous configuration tasks.

To upload the user and group data, perform the following steps:

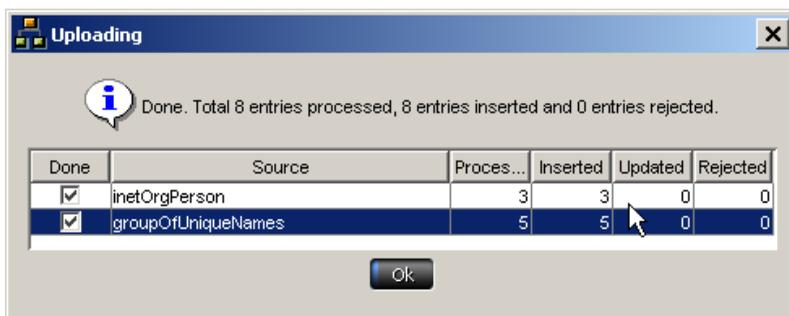
1. On the RadiantOne Synchronization Services Administration console, in the right pane, click the Deployment tab, click the folder icon, and open the topology that you just saved. The topology is displayed.
2. From the Deployment drop-down menu, select Upload. The upload topology page is displayed (see [Figure 6-16](#)).

Figure 6-16 Upload Topology Page



3. Check on both Delete Rows check boxes and click Test Upload. The Uploading page is displayed and indicates that the upload is successful (see [Figure 6-17](#)).

Figure 6-17 Topology Uploading Page



4. Click Ok
5. On the upload topology page, click Actual Upload (see Figure 6-16). The Uploading page is displayed again and indicates whether all entries were processed successfully.
6. Click Ok and on the Upload - Topology page, click close.
7. To verify that the upload actually moved user data into the designated WebLogic Enterprise Security identity directory, perform the following steps:
 - a. Go to the WebLogic Enterprise Security Administration Console and, in the left pane, open the Identity folder.
 - b. Click Groups and Users. The user data is displayed in the console.
8. This completes the user and group data upload.

Configuring Metadirectory Synchronization

This section describes how to configure the metadirectory components for automatic updates to the policy database whenever changes are made to the user repository.

The RadiantOne Synchronization Services provides connectors and a synchronization hub that work together to synchronize data between various data sources. Connectors interface with the data sources. Data flows to and from the connectors asynchronously in the form of XML messages. All messages flow through the synchronization hub, which is a server that transforms the messages and routes them to the connectors that are subscribed to the changes. The BEA WebLogic JMS messaging broker manages the topics and provides guaranteed message delivery.

The role of the connectors is two fold. First, the connectors capture changes in the data source, translate the changes into a common XML format, and send them to the synchronization hub through the messaging server. Secondly, the connectors receive XML messages, translate them, and apply the changes to the data source.

To configure metadirectory synchronization, perform the following tasks:

- [“Configuring the Synchronization Hub” on page 6-30](#)
- [“Configuring the Directory Connector” on page 6-30](#)
- [“Configuring the Policy Database Connectors” on page 6-31](#)
- [“Starting the Synchronization Hub” on page 6-33](#)
- [“Starting the Source and Destination Connectors” on page 6-33](#)

Configuring the Synchronization Hub

To configure the synchronization hub, set the parameters in the `RadiantOne_Home\r1syncsvcs\rlicon.ini` file to match the settings on the WebLogic Server. The required settings are shown in [Listing 6-1](#). This information is used by the RadiantOne Synchronization Services tool to contact the JMS server running on the WebLogic Enterprise Security Administration Server.

Listing 6-1 Synchronization Hub Settings

```
...
[BEA]
JMS_SERVER_NAME=RLI_JMS_SERVER
JNDI_CLASS_NAME=weblogic.jndi.WLInitialContextFactory
CONNECTION_FACTORY_NAME=weblogic.asiAdminServer.jms.TopicConnectionFactory
URL=t3://localhost:7000
USER_NAME=system
USER_PASSWORD=weblogic
Msg Time-To-Live=3600000
RETRY_PERIOD=12
MAXIMUM_ATTEMPTS = 11
```

Configuring the Directory Connector

Note: This task is necessary only if you are using the Sun ONE Directory Server for your user repository. If you are using another type of user repository server, such as Active Directory, Windows NT, or a database, skip this section and go [“Configuring the Policy Database Connectors” on page 6-31](#). For more information on LDAP connector configuration requirements and procedures, see the *RadiantOne Synchronization Services Guide* located in the RadiantOne installation directory.

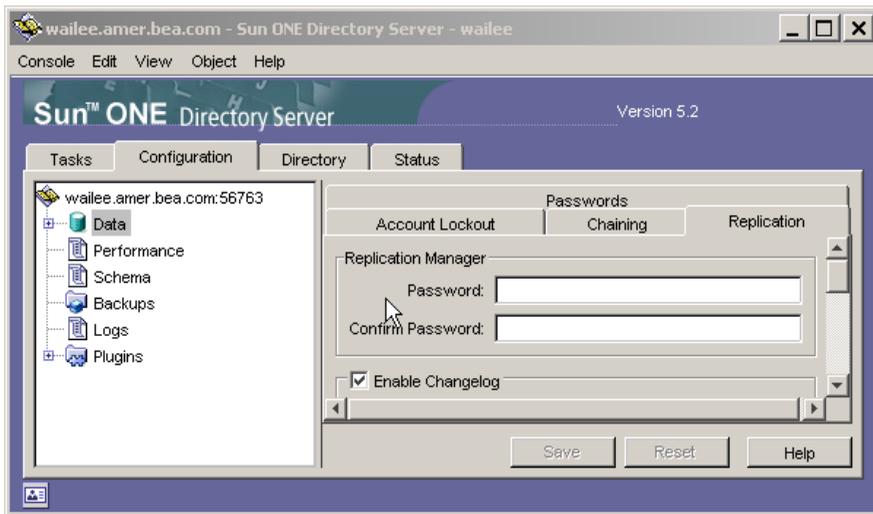
If you are using the Sun ONE Directory Server, you must configure the directory connector for so that changes to the user repository are automatically updated in the policy database.

To configure the directory connector, perform the following steps:

1. Click `Start>Programs>Sun ONE Server Products>Sun ONE Server Console 5.2` and log into the Directory Server Console using the `admin` User ID.

2. In the left pane, expand the Domain and Server Group nodes, and double-click the Directory Server for your directory server. The Directory Server Tasks page is displayed.
3. Click the Configuration tab, select the Data node in the left pane, and select the Replication tab. The Replication page is displayed (see [Figure 6-18](#)).

Figure 6-18 Directory Server Replication Page



4. Check the Enable Changelog check box and click Save.
 5. In the left pane, open the Plugins folder, click Retro Changelog Plugin, check the Enable plug check box, and click Save.
- Note:** If you are using an LDAP server other than iPlanet or you do not want to use the iPlanet changelog, Radiant Synchronization Services also has a polling connector. For information on configuring the other LDAP connectors, refer to the RadiantOne online documentation located in the RadiantOne installation directory.
6. Click the Tasks tab and click Restart Directory Server to restart the server.
 7. The completes configuration of the directory connector.

Configuring the Policy Database Connectors

To configure the policy database connector, perform the following steps:

1. On the RadiantOne Synchronization Services Administration console, click the topology tab and open the topology.
2. Right-click the ASI_USERS database object, and select Configure. The ASI_USERS Destination Configuration page is displayed (see [Figure 6-19](#)).

Figure 6-19 ASI_USERS Destination Configuration Page



- Note:** If the database object that the connector is listening to changes or is modified (for example, one of the data types changes or you add or remove columns), you can reconfigure the connector by right-clicking on the database object in the topology, and then choosing Configure. A note is displayed on the Login Info tab that specifies how the connector is configured. Enter the user and password information and, on the Initialization tab, reconfigure the connector by either Applying the Script or Saving and executing it later.
3. Enter the user and password to connect to the database as the database administrator and click Connect. A “Connection Successful” dialog is displayed. Click Ok. A script is generated to create the rli_con user, the needed log tables, and triggers.
 4. Click the Initialization tab, select the Apply Now radio button to generate log tables and triggers for the ASI_USERS database object, and click Apply. A “Configuration completed” dialog box is displayed. Click Ok.
 5. Click Ok.
 6. Repeat steps 2 through 5 for the ASI_GROUPS database object.
 7. This completes the configuration of the policy database connectors.

Starting the Synchronization Hub

To start the Synchronization Hub, perform the following steps:

1. On the RadiantOne Synchronization Services Administration console, click the Deployment drop-down menu, select Synchronization Hub, and click Start. The JMS Connection Username/Password dialog is displayed.
2. Enter the JMS connection username and password (if necessary) and click Ok. Use the same username/password that you used to log into the WebLogic Server Administration Console. A small window for the hub opens and indicates that the hub is running.

Note: You can also use the Start Hub icon on the Deployment Tab to start the hub.

Starting the Source and Destination Connectors

To start the source and destination connectors, perform the following steps:

Note: The Synchronization Hub must be running before you start the connectors.

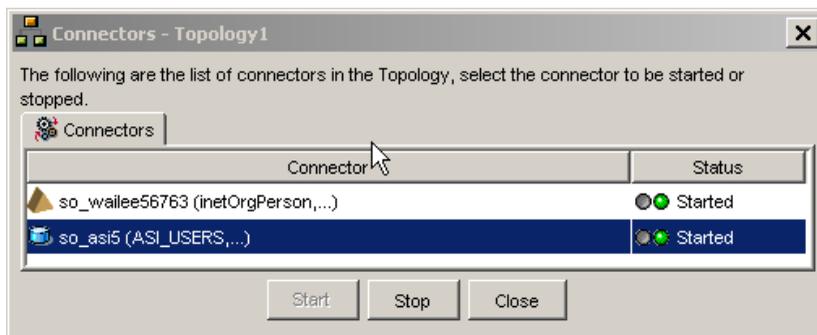
To start the Synchronization Hub, perform the following steps:

1. On the RadiantOne Synchronization Services Administration console, select the Deployment tab, click the folder icon, and select and open the Topology. The topology is displayed.

Note: You cannot modify the topology when you open it using the Deployment tab. To modify a topology you must open it using the Topology tab.

2. From the Deployment drop-down menu, and click Connectors. The Connectors - Topology dialog is displayed (see [Figure 6-20](#)).

Figure 6-20 Connectors - Topology Page



3. Select each connector and click Start. The connectors start. A small window for each connector opens and indicates that the connector is running.

Note: The connectors can also be started and stopped by clicking the Start Connector and Stop Connectors icons under the Deployment tab.

4. This completes configuration of metadirectory Synchronization.

Verifying that Metadirectory Synchronization Works

This section describes how to verify that metadirectory synchronization is properly configured such that changes to user and group entries in the user repository are reflected in the policy database.

To verify that metadirectory synchronization is properly configured, perform the following steps:

1. Use the user repository server to create a new user and add that user to the source group.
2. Open the Administration Console, open the Identity folder, the list of identity directories is displayed in the right pane.
3. Select the identity directory that you configured for automatic updates (for example, `ldapdir`), and click Users in the left pane. The new user is displayed in the list of users in the right pane.

Uninstalling

The following sections describe how to uninstall the Administration Application from both Windows and UNIX platforms:

Note: If you have entered security policy and configuration information the Administration Console and you want to save it, you must export it, uninstall the Administration Application, re-install the Administration Application, and import the security policy and configuration information. For instructions on exporting and importing policy and configuration information, see “[Exporting Policy Data](#)” and “[Importing Policy Data](#)” in the *Policy Managers Guide*.

- [“Uninstalling the Administration Application on Windows” on page 7-1](#)
- [“Uninstalling the Administration Application on Solaris or Linux” on page 7-3](#)

Uninstalling the Administration Application on Windows

This procedure removes the Administration Application, which includes the Administration Console and all scripts that relate to management tasks. It does not remove any users and groups that are installed during the installation procedure; you need to remove these manually after the uninstall completes.

To uninstall the Administration Application, do the following:

1. Shut down any servers and services that are running.
2. Click Start and select Programs>BEA WebLogic Enterprise Security>Uninstall Administration Application.

The Uninstall Welcome window appears.

3. Click Next.

The BEA Uninstaller - Administration Application window appears and the uninstall process begins.

As the uninstall process runs, a checklist is displayed, listing the uninstallation tasks as they complete. After the Administration Application is removed, the "uninstall complete" message appears.

Note: Make sure to check the Uninstall SCM box also. If you do not do this, you will have to delete these product files manually.

4. Click Done.

You have successfully removed Administration Application from your computer.

After the uninstall completes, you may notice that the product directory and various log files remain in the product directory. You may remove these manually or you may want to keep them. This procedure does not remove the WebLogic Server software. You must remove that software according to WebLogic Server documentation. For additional instructions for completely removing the product, see [“Additional Steps” on page 7-2](#).

Additional Steps

After the uninstall completes, you may notice that the product directory and various log files remain in the product directory. You may remove these manually or you may want to keep them. This procedure does not remove the WebLogic Server software. You must remove this software according to WebLogic Server documentation.

1. Uninstall all WebLogic Enterprise Security products (as described in [“Uninstalling the Administration Application on Windows” on page 7-1](#)), including the Service Control Manager (SCM). This removes these product references from the Windows registry but does not remove all files from the product directories. These files are preserved so that audit and log information is not lost.
2. Delete the remainder of the product directories. Doing this removes all instance files that remain, unless you created them outside of the product directory. In this case, you need to go to the directory where you installed the instance and delete the instance directory.
3. Delete shortcuts from the start menu. You must delete product Programs menu shortcuts manually. Right click on each shortcut you want to delete, and then choose delete.

Note: If you do not delete the menu shortcuts and then reinstall the product, duplicate product names appear in the Programs menu.

4. To delete users and groups, open the Control Panel>Administrative Tools>Computer Management>Local Users and Groups window.
5. Delete the following users and groups:
 - The Administration Application user (asiadmin by default)
 - The Service Control Manager user (scmuser by default)
 - The WLES administrators group (asiadgrp)
 - The WLES users group (asiusers)

If you know the passwords of asiadmin and scmuser (passwords typed in during a previous install, rather than accepting the defaults), then you may leave these users in place and enter those passwords if you want to reinstall the product.
6. To remove the ActiveX controls, open the Control Panel>Add or Remove Programs and uninstall the ActiveX control named "BEA WLES Administration Console". You need to do this on each Windows system previously used to connect to an Administration Server.

Uninstalling the Administration Application on Solaris or Linux

To run the graphical mode uninstallation program, your console must support a Java-based GUI. If the uninstallation program determines that your system cannot support a Java-based GUI, the uninstallation program automatically starts in console mode.

To uninstall the Administration Application software:

1. Log in to the machine as root (or su root).
2. Shut down any servers and services that are running.
3. Open a command shell and go to the directory where you installed the product, for example:

```
BEA_HOME/wles42-admin/uninstall
```

where *BEA_HOME* represents the directory in which you installed product.

4. At the command prompt, type `uninstall.sh`.

The BEA Uninstaller - Administration Application window appears and the uninstall process begins.

Uninstalling

Note: If your system supports a graphical user interface, the uninstall program starts in graphical mode. If your system does not support a graphical user interface, the uninstall program starts in console mode.

5. Respond to the prompts to uninstall the product.

Index

D

- database backup
 - Oracle 3-31
 - Sybase 3-55
- database schema
 - installing, Linux 5-5
 - installing, Sun Solaris 5-4
 - installing, Windows 5-3
- distribution
 - CD-ROM 2-2
 - product 2-1

E

- environment variable 5-2

I

- Installation 2-1
- installation 2-1, 2-6, 5-2

M

- metadirectory synchronization service
 - installing 6-5

P

- password 4-4
 - system, changing 5-7
- policy database
 - administration 3-31
 - Sybase 3-55
- processes

starting 5-6

S

- Sybase Adaptive Server
 - installing 3-35
- system
 - security infrastructure 4-2
- system requirements 2-1

U

- user names 4-4

