



BEA WebLogic Enterprise

Getting Started

WebLogic Enterprise 5.1
Document Edition 5.1
May 2000

Copyright

Copyright © 2000 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, ObjectBroker, TOP END, and Tuxedo are registered trademarks of BEA Systems, Inc. BEA Builder, BEA Connect, BEA Manager, BEA MessageQ, BEA Jolt, M3, eSolutions, eLink, WebLogic, and WebLogic Enterprise are trademarks of BEA Systems, Inc.

All other company names may be trademarks of the respective companies with which they are associated.

Getting Started

Document Edition	Part Number	Date	Software Version
5.1	861-001001-004	May 2000	BEA WebLogic Enterprise 5.1

Contents

About This Document

What You Need to Know	viii
e-docs Web Site	viii
How to Print the Document.....	ix
Related Information.....	ix
Contact Us!	ix
Documentation Conventions	x

Part I. Overview of the WebLogic Enterprise Product and Programming Environments

1. Overview of the WebLogic Enterprise Product

Product Overview	1-1
J-Engine Features	1-3
T-Engine Features	1-4

2. The WebLogic Enterprise CORBA Programming Environment

Overview of the WebLogic Enterprise CORBA Programming Features	2-1
IDL Compilers.....	2-2
Development Commands	2-3
Administration Tools.....	2-4
ActiveX Application Builder	2-6
WebLogic Enterprise CORBA Object Services.....	2-8
WebLogic Enterprise Architectural Components	2-9
Bootstrap Object.....	2-11
IIOP Listener/Handler	2-12
ORB.....	2-13

TP Framework	2-14
How WebLogic Enterprise CORBA Client and Server Applications Interact	2-16
Step 1: The Server Application Is Initialized	2-17
Step 2: The Client Application Is Initialized.....	2-18
Step 3: The Client Application Authenticates Itself to the WebLogic Enterprise Domain.....	2-18
Step 4: The Client Application Obtains a Reference to the Object Needed to Execute Its Business Logic	2-19
Step 5: The Client Application Invokes an Operation on the CORBA Object.....	2-21

3. The WebLogic Enterprise JavaBeans (EJB) Programming Environment

Overview of the WebLogic Enterprise EJB Programming Environment	3-2
Types of Beans Supported in WebLogic Enterprise.....	3-3
EJBs and Persistence	3-5
Roles of People Who Develop, Build, Deploy, and Administer EJBs	3-5
Items You Create for an EJB Application	3-9
Tools and Facilities Provided for Building and Deploying EJBs.....	3-10
EJBs and Failover in the WebLogic Enterprise Environment.....	3-11

Part II. Developing WebLogic Enterprise CORBA Applications

4. Developing WebLogic Enterprise CORBA Applications

Overview of the Development Process for WebLogic Enterprise CORBA Applications.....	4-2
The Simpapp Sample Application	4-4
Step 1: Write the OMG IDL Code.....	4-5
Step 2: Generate Client Stubs and Skeletons.....	4-6
Step 3: Write the Server Application.....	4-8
Writing the Methods That Implement Each Interface's Operations	4-9
Creating the Server Object	4-12
Defining an Object's Activation Policies	4-16
Creating and Registering a Factory	4-17
Releasing the Server Application	4-19
Step 4: Write the Client Application	4-21

Step 5: Create an XA Resource Manager.....	4-24
Step 6: Create a Configuration File	4-25
Step 7: Create the TUXCONFIG File	4-28
Step 8: Compile the Server Application.....	4-29
Step 9: Compile the Client Application.....	4-30
Step 10: Start the WebLogic Enterprise CORBA Application	4-30
Additional WebLogic Enterprise CORBA Sample Applications	4-31

5. Using Security

Overview of the Security Service	5-1
How Security Works	5-2
The Security Sample Application.....	5-4
Development Steps.....	5-6
Step 1: Define the Security Level in the Configuration File.....	5-6
Step 2: Write the CORBA Client Application	5-7

6. Using Transactions

Overview of the Transaction Service	6-1
What Happens During a Transaction	6-3
Transactions Sample Application.....	6-4
Development Steps.....	6-7
Step 1: Write the OMG IDL Code	6-8
Step 2: Define Transaction Policies for the Interfaces	6-10
Step 3: Write the CORBA Client Application	6-11
Step 4: Write the Server Application	6-13
Step 5: Create a Configuration File.....	6-14

Part III. Developing WebLogic Enterprise EJB Applications

7. Developing WebLogic Enterprise EJB Applications

Overview of the Development Process for WebLogic Enterprise EJB Applications.....	7-3
The statefulSession EJB Sample Application	7-5
Developing EJB Applications	7-7
Step 1: Create the EJB.....	7-8
Step 2: Create the Module Initializer Object.....	7-13

Step 3: Create the Deployment Descriptor.....	7-14
Step 4: Create a Standard EJB JAR File	7-19
Building and Deploying EJB Applications	7-20
Step 5: Create the WebLogic EJB Extensions to the Deployment Descriptor DTD	7-21
Step 6: Modify the Deployment Descriptor	7-26
Step 7: Package the Components Into a Deployable EJB JAR File	7-28
Step 8: Configure the EJB Application	7-29
Step 9: Create the Client Application	7-33
Step 10: Start and Run the WebLogic Enterprise EJB Application	7-36
Step 11: Dynamically Manage the EJB Deployment (Hot Redeployment)	7-37
WebLogic Enterprise EJB Sample Applications	7-39

8. Designing Enterprise JavaBeans for the WebLogic Enterprise System

Designing EJB Applications for the WebLogic Enterprise System	8-2
The Client Application Programmer's View.....	8-2
The EJB Programmer's View.....	8-5
EJBs and Persistence	8-11
Development Considerations for EJBs and Persistence.....	8-11
Container-managed Entity Beans	8-12
Bean-managed Entity Beans.....	8-14
Stateful Session Beans.....	8-18
Stateless Session Beans	8-21

Index

About This Document

This document presents an overview of the BEA WebLogic Enterprise™ product and describes the development process for developing distributed CORBA and Enterprise JavaBeans (EJB) applications using the WebLogic Enterprise software.

This document does not discuss every feature of the WebLogic Enterprise product; instead, it gives a general description of building a typical application or bean using the WebLogic Enterprise programming environment. To find information about all the WebLogic Enterprise features, see the home page of the WebLogic Enterprise online documentation.

This document covers the following topics:

- Chapter 1, “Overview of the WebLogic Enterprise Product,” presents an overview of the WebLogic Enterprise product.
- Chapter 2, “The WebLogic Enterprise CORBA Programming Environment,” describes the CORBA programming environment available in the WebLogic Enterprise product and the architectural components of the CORBA programming environment.
- Chapter 3, “The WebLogic Enterprise JavaBeans (EJB) Programming Environment,” describes the EJB programming environment available in the WebLogic Enterprise product and the architectural components of the EJB programming environment.
- Chapter 4, “Developing WebLogic Enterprise CORBA Applications,” explains how to build a typical WebLogic Enterprise CORBA application, using the Simpapp sample application as an example.
- Chapter 5, “Using Security,” describes how security is incorporated into a WebLogic Enterprise CORBA application. The Security sample application is used as an example.

- Chapter 6, “Using Transactions,” describes how transactions are incorporated into a WebLogic Enterprise CORBA application. The Transactions sample application is used as an example.
- Chapter 7, “Developing WebLogic Enterprise EJB Applications,” explains how to build a typical WebLogic Enterprise EJB application, using the `statefulSession` EJB application, which is shipped with the WebLogic Enterprise software, as an example.
- Chapter 8, “Designing Enterprise JavaBeans for the WebLogic Enterprise System,” explains how to design a typical EJB application using the WebLogic Enterprise programming environment.

What You Need to Know

This document is intended for programmers who want to familiarize themselves with the WebLogic Enterprise programming environment and create either distributed CORBA or Enterprise JavaBeans applications using the WebLogic Enterprise product.

e-docs Web Site

The BEA Weblogic Enterprise product documentation is available from the BEA Systems, Inc. corporate Web site. From the BEA Home page, click the Product Documentation button or go directly to the “e-docs” Product Documentation page at <http://e-docs.bea.com>.

How to Print the Document

You can print a copy of this document from a Web browser, one file at a time, by using the File—>Print option on your Web browser.

A PDF version of this document is available on the WebLogic Enterprise documentation Home page on the e-docs Web site (and also on the documentation CD). You can open the PDF in Adobe Acrobat Reader and print the entire document (or a portion of it) in book format. To access the PDFs, open the WebLogic Enterprise documentation Home page, click the PDF Files button, and select the document you want to print.

If you do not have the Adobe Acrobat Reader installed, you can download it for free from the Adobe Web site at <http://www.adobe.com/>.

Related Information

For more information about CORBA, Java 2 Enterprise Edition (J2EE), BEA Tuxedo®, distributed object computing, transaction processing, C++ programming, and Java programming, see the *Bibliography* in the WebLogic Enterprise online documentation.

Contact Us!

Your feedback on the BEA WebLogic Enterprise documentation is important to us. Send us e-mail at docsupport@bea.com if you have questions or comments. Your comments will be reviewed directly by the BEA professionals who create and update the WebLogic Enterprise documentation.

In your e-mail message, please indicate that you are using the documentation for the BEA WebLogic Enterprise 5.1 release.

If you have any questions about this version of BEA WebLogic Enterprise, or if you have problems installing and running BEA WebLogic Enterprise, contact BEA Customer Support through BEA WebSUPPORT at www.bea.com. You can also contact Customer Support by using the contact information provided on the Customer Support Card, which is included in the product package.

When contacting Customer Support, be prepared to provide the following information:

- Your name, e-mail address, phone number, and fax number
- Your company name and company address
- Your machine type and authorization codes
- The name and version of the product you are using
- A description of the problem and the content of pertinent error messages

Documentation Conventions

The following documentation conventions are used throughout this document.

Convention	Item
boldface text	Indicates terms defined in the glossary.
Ctrl+Tab	Indicates that you must press two or more keys simultaneously.
<i>italics</i>	Indicates emphasis or book titles.

Convention	Item
monospace text	Indicates code samples, commands and their options, data structures and their members, data types, directories, and filenames and their extensions. Monospace text also indicates text that you must enter from the keyboard. <i>Examples:</i> <code>#include <iostream.h> void main () the pointer psz chmod u+w *\br/>\tux\data\ap .doc tux.doc BITMAP float</code>
monospace boldface text	Identifies significant words in code. <i>Example:</i> <code>void commit ()</code>
<i>monospace</i> <i>italic</i> <i>text</i>	Identifies variables in code. <i>Example:</i> <code>String <i>expr</i></code>
UPPERCASE TEXT	Indicates device names, environment variables, and logical operators. <i>Examples:</i> <code>LPT1 SIGNON OR</code>
{ }	Indicates a set of choices in a syntax line. The braces themselves should never be typed.
[]	Indicates optional items in a syntax line. The brackets themselves should never be typed. <i>Example:</i> <code>buildobjclient [-v] [-o name] [-f <i>file-list</i>]... [-l <i>file-list</i>]...</code>
	Separates mutually exclusive choices in a syntax line. The symbol itself should never be typed.

Convention	Item
...	<p>Indicates one of the following in a command line:</p> <ul style="list-style-type: none">■ That an argument can be repeated several times in a command line■ That the statement omits additional optional arguments■ That you can enter additional parameters, values, or other information <p>The ellipsis itself should never be typed.</p>
	<p><i>Example:</i></p> <pre>buildobjclient [-v] [-o name] [-f <i>file-list</i>]... [-l <i>file-list</i>]...</pre>
.	<p>Indicates the omission of items from a code example or from a syntax line.</p>
.	<p>The vertical ellipsis itself should never be typed.</p>
.	

Part I Overview of the WebLogic Enterprise Product and Programming Environments

- Chapter 1. Overview of the WebLogic Enterprise Product
- Chapter 2. The WebLogic Enterprise CORBA Programming Environment
- Chapter 3. The WebLogic Enterprise JavaBeans (EJB) Programming Environment

1 Overview of the WebLogic Enterprise Product

This topic includes the following sections:

- Product Overview
- J-Engine Features
- T-Engine Features

Product Overview

The BEA WebLogic Enterprise product features the integration of BEA Systems Inc. industry-leading technologies — a high performance Java application server and a scalable back-end transaction platform. These technologies are now merged into a single product that gives you maximum choice and flexibility in building robust e-commerce applications that extend from the Web to the enterprise.

The J-Engine in the WebLogic Enterprise product is built upon the BEA WebLogic Server™ technology. The J-Engine offers an award-winning Java application server that incorporates the most comprehensive implementation of the Java 2 Enterprise Edition (J2EE) standards. The application server provides the foundation for the rapid development of Web applications and the performance and reliability required for mission-critical e-commerce sites.


1 Overview of the WebLogic Enterprise Product

The T-Engine in the WebLogic Enterprise product is built upon the proven infrastructure provided in the BEA Tuxedo™ product. The T-Engine delivers a scalable transaction platform with unparalleled choice of development and deployment options. Using the T-Engine, you can build integrated enterprise applications using multiple programming models. CORBA, J2EE, and Tuxedo applications can all be developed with fully integrated transaction management, security, administration, and reliability capabilities.

The connection pooling technology incorporated into the WebLogic Enterprise product provides for scalable connectivity between the J-Engine and T-Engine environments.

Figure 1-1 illustrates the WebLogic Enterprise product.

Figure 1-1 The WebLogic Enterprise Product



The following sections outline the features of the T-Engine and J-Engine in the WebLogic Enterprise product.

J-Engine Features

The J-Engine in WebLogic Enterprise product provides the following set of features:

- A complete implementation of EJB, servlets, Java Server Pages (JSPs), Java Message Service (JMS), Java Database Connectivity (JDBC), and the Java Naming and Directory Interface (JNDI) as specified by the J2EE standard.
- Web page and component clustering of EJBs across multiple servers.
Web page clustering handles transparent replication, load-balancing, and failover for presentation logic. Component clustering handles the complexities of replication, load-balancing, and failover for EJBs.
- JSPs for easy development and deployment of dynamic Web content. JSPs can be used with personalization, database access, and transaction EJBs to develop many kinds of high-performance Web applications.
- Remote Method Invocation (RMI). With RMI, an application can use distributed objects as easily as local objects. RMI can be clustered across multiple WebLogic Servers.
- JMS for applications requiring real-time information about changing application conditions. The JMS implementation in the J-Engine provides store and forward and point-to-point messaging with guaranteed delivery. JMS also provides a publish/subscribe event management model.
- Multitiered JDBC which allows a Java application to access and update databases from anywhere on the network. The J-Engine includes its own native JDBC drivers for leading database products and also works with any third-party JDBC driver.
- Support for Wireless Markup Language (WML) and integration with leading WAP servers for support of wireless clients.
- XML support for any XML-compliant browser.

1 Overview of the WebLogic Enterprise Product

- Interoperability with Microsoft COM objects. Microsoft COM objects can be integrated into the Weblogic Server environment, wrapped with a Java class, and transparently shared over the network.
- Access for Web browsers directly via HTTP request. Forwarding capabilities, such as HTTP proxy support, enable dispatching to servers other than the original Web server.
- Support for the Java Transaction application programming interface (API) which allows client or server applications to initiate transactions that are propagated to other servers.
- A Java management console for remotely monitoring and updating the state of applications and clusters. SNMP support which allows use of any third-party, SNMP-compliant management framework.
- A Zero Administration Client (ZAC) which supports the automatic distribution of Java applets, applications, or systems. With ZAC, program libraries, even a new WebLogic Server release, can be installed centrally by an administrator.
- Dynamic application partitioning and cluster membership.
- Support for Oracle, Informix, Sysbase, and MS SQL-server databases.
- Integration with Integrated Development Environments (IDEs) including IBM VisualAge, Inprise Jbuilder, Microsoft Visual J++ or any Java 1.1 or higher compliant IDE.
- Integrated security and firewall support. Network applications are secured with optional encryption, authentication, and authorization based on the SSL protocol, X.509 digital certificates, and access control lists (ACLs).
- Management of IIOP connections that allows the restarting of a connection pool without affecting the availability of the Web environment.

T-Engine Features

The T-Engine in WebLogic Enterprise product provides the following set of features:

- A full suite of server-side components including:

- BEA Tuxedo
- A CORBA C++ ORB
- A CORBA Java ORB
- EJB container
- RMI support

One or more BEA Tuxedo, CORBA Java, CORBA C++, EJB, or RMI server components can be deployed in a single WebLogic Enterprise application.

The WebLogic Enterprise product supports version 1.1 of the Sun Microsystems, Inc. Enterprise JavaBeans Specification.

- Rich clients options including:
 - Tuxedo /workstation client
 - A CORBA C++ ORB client
 - A CORBA Java ORB client
 - A WebLogic Enterprise RMI client
 - An ActiveX client
- A proven runtime infrastructure for hosting e-commerce transaction applications, including client connection concentrators, high-performance message routing and load balancing, high-availability features, and database connection pooling.
- Full support for Tuxedo 6.x applications.
- EJB version 1.1 support for Enterprise Java Bean applications.
- A Transaction Processing (TP) Framework for object state and transaction management in CORBA applications.
- Interoperability with IIOP-compliant ORBs such as the JDK 1.2 Java ORB.
- Access to databases from Java applications using two-phase commit via BEA JDBC-XA drivers. Support for XA-compliant databases using Tuxedo or C++. Drivers are provided for Oracle 8.0.5 and 8.1.5 databases.
- A Management Information Base (MIB) that defines the key management attributes of WebLogic Enterprise applications. In addition, programming interfaces and scripting capabilities are available to access the MIBs.

1 Overview of the WebLogic Enterprise Product

- An Administration Console graphical user interface (GUI) for the management of the WebLogic Enterprise environment.
- Hot deployment of EJBs through the Deployer GUI tool.
- The Java Naming and Directory Interface (JNDI) used by WebLogic Enterprise client applications to find WebLogic Enterprise server-side EJB Home objects and RMI objects.
- The CORBA and Java Transaction Services (OTS and JTS) to ensure the integrity of your data even when transactions span multiple programming models, databases, and applications.
- A security service that handles authentication for principals that need to access resources in a CORBA object or EJB in the WebLogic Enterprise environment. Access control lists (ACLs) are also provided for EJBs in your WebLogic Enterprise application.
- The Secure Sockets Layer (SSL) protocol to encrypt client to server communication on the wire. SSL support includes IIOP connection pools between the J-Engine and the T-Engine.
- Propagation of the security context from the J-Engine to the T-Engine through IIOP connection pools.
- A Security Service Plug-In Interface (SPI) for CORBA that allows integration of third-party security plug-ins.
- A Notification Service that receives event posting messages, filters them, and distributes the messages to subscribers. The Notification Service provides two sets of interfaces: a CORBA-based interface and a simplified BEA-proprietary interface.
- An implementation of the CosLifeCycle service.
- An implementation of CosNaming that allows WebLogic Enterprise CORBA server applications to advertise object references using logical names.
- An interface repository that stores meta information about WebLogic Enterprise CORBA objects. Meta information includes information about modules, interfaces, operations, attributes, and exceptions.
- Dynamic Invocation Interface (DII) support. DII allows WebLogic Enterprise CORBA client applications to dynamically create requests for objects that were not defined at compile time.

- Jolt for client-side access to BEA Tuxedo services. Jolt enables browser-based clients (both Java applets and applications) to invoke BEA Tuxedo services and process the results. Jolt is installed separately from the T-Engine.
- Java Enterprise Tuxedo (JET) application programming interface (API) for server-side access to BEA Tuxedo services. The JET API enables Java server applications (CORBA Java, EJB, or RMI) running within a WebLogic Enterprise domain to invoke BEA Tuxedo services and process the results. The JET API is automatically installed when you install the WebLogic Enterprise product.

The rest of this *Getting Started* manual describes the programming environment of the T-Engine and the development process for CORBA objects and EJBs in the T-Engine environment. For a description of the programming environment and development process for the J-Engine, see the *Overviews* topic in the Weblogic Server portion of the online documentation CD.

1 *Overview of the WebLogic Enterprise Product*

2 The WebLogic Enterprise CORBA Programming Environment

This topic includes the following sections:

- Overview of the WebLogic Enterprise CORBA Programming Features
- WebLogic Enterprise CORBA Object Services
- WebLogic Enterprise Architectural Components
- How WebLogic Enterprise CORBA Client and Server Applications Interact

Overview of the WebLogic Enterprise CORBA Programming Features

The WebLogic Enterprise product offers a robust CORBA programming environment that simplifies the development and management of distributed objects. The following topics describe the features of the programming environment:

- IDL Compilers
- Development Commands
- Administration Tools
- ActiveX Application Builder

IDL Compilers

The WebLogic Enterprise product comes with two IDL compilers that make object development easier:

- `idl`—compiles the OMG IDL file and generates client stub and server skeleton files required for interface definitions being implemented in C++.
- `idltojava`—compiles IDL files to Java source code based on IDL-to-Java mappings defined by the OMG. The `idltojava` compiler provided with the WebLogic Enterprise product includes several enhancements, extensions and additions that are not present in the original Sun Microsystems, Inc. version of the compiler. The WebLogic Enterprise specific revisions are summarized below.
 - Differs from that described in the Sun Microsystems, Inc. documentation in behavior and defaults of the flags.
 - Includes a new `#pragma` tag: `#pragma ID <name> <Repository_id>`.
 - Includes a new `#pragma` tag: `#pragma version <name> <m..n>`.
 - Extends the `#pragma` prefix to work on inner scope. A blank prefix reverts.
 - Allows unions with boolean discriminators.
 - Allow declarations nested inside complex types.
- `m3idltojava`—compiles the OMG IDL file and generates client stub and server skeleton files required for interface definitions being implemented in Java.

For a description of how to use the IDL compilers, see Chapter 4, “Developing WebLogic Enterprise CORBA Applications.”

For a description of the `idl`, `idltojava`, and `m3idltojava` commands, see [*Commands, Server Processes, and MIB Reference*](#) in the WebLogic Enterprise online documentation.

Development Commands

Table 2-1 lists the commands that the WebLogic Enterprise product provides for developing CORBA application components and managing the Interface Repository.

Table 2-1 WebLogic Enterprise CORBA Development Commands

Development Command	Description
<code>buildjavaserver</code>	Constructs a server application JAR file for a Java server application.
<code>buildobjclient</code>	Constructs a C++ client application.
<code>buildobjserver</code>	Constructs a C++ server application.
<code>buildXAJS</code>	Constructs an XA resource manager to be used with a Java server application group.
<code>genicf</code>	Generates an Implementation Configuration File (ICF). The ICF file defines activation and transaction policies for C++ server applications.
<code>idl2ir</code>	Creates the Interface Repository and loads interface definitions into it.
<code>ir2idl</code>	Shows the content of the Interface Repository.
<code>irdel</code>	Deletes the specified object from the Interface Repository.

For a description of how to use the development commands to develop client and server applications, see Chapter 4, “Developing WebLogic Enterprise CORBA Applications.”

For a description of the development commands, see [*Commands, Server Processes, and MIB Reference*](#) in the WebLogic Enterprise online documentation.

Administration Tools

The WebLogic Enterprise product provides a complete set of tools for administering your WebLogic Enterprise environment. You can manage the WebLogic Enterprise application through commands, through a graphical user interface, or by including administration utilities in a script.


You can use the commands listed in Table 2-2 to perform administration tasks for your WebLogic Enterprise application.

Table 2-2 WebLogic Enterprise Administration Commands

Administration Command	Description
tmadmin	Displays information about current configuration parameters.
tmboot	Activates the WebLogic Enterprise application referenced in the specified configuration file. Depending on the options used, the entire application or parts of the application are started.
tmconfig	Dynamically updates and retrieves information about the configuration of a WebLogic Enterprise application.
tmloadcf	Parses the configuration file and loads the binary version of the configuration file.
tmshutdown	Shuts down a set of specified server applications, or removes interfaces from a configuration file.
tmunloadcf	Unloads the configuration file.

The Administration Console is a Java-based applet that you can download into your Internet browser and use to remotely manage your WebLogic Enterprise applications. The Administration Console allows you to perform administration tasks, such as monitoring system events, managing system resources, creating and configuring administration objects, and viewing system statistics. Figure 2-1 shows the main window of the Administration Console.

Figure 2-1 Administration Console Main Window



In addition, a set of utilities called the AdminAPI is provided for directly accessing and manipulating system settings in the Management Information Bases (MIBs) for the WebLogic Enterprise product. The advantage of the AdminAPI is that it can be used to automate administrative tasks, such as monitoring log files and dynamically reconfiguring an application, thus eliminating the need for manual intervention.

For information about the Administration commands, see [Commands, Server Processes, and MIB Reference](#) and [Administration](#) in the WebLogic Enterprise online documentation.

For a description of the Administration Console and how it works, see the online help that is integrated into the Administration Console graphical user interface (GUI).


For information about the Admin API, see [BEA Tuxedo Reference](#) in the WebLogic Enterprise online documentation.

ActiveX Application Builder

The ActiveX Application Builder is a development tool that you use with a client development tool (such as Visual Basic) to select which CORBA interfaces in a WebLogic Enterprise domain you want your ActiveX client application to interact with. In addition, you use the ActiveX Application Builder to create Automation bindings for CORBA interfaces, and to create packages for deploying ActiveX views of CORBA objects to client machines.

Figure 2-2 shows the ActiveX Application Builder main window.

Figure 2-2 ActiveX Application Builder Main Window



For a description of the ActiveX Application Builder and how it works, see the online help that is integrated into the ActiveX Application Builder graphical user interface (GUI). For information about creating ActiveX client applications, see the PDF version of the *WebLogic Enterprise ActiveX Client Developer's Guide* in the WebLogic Enterprise online documentation.

WebLogic Enterprise CORBA Object Services

The WebLogic Enterprise product includes a set of environmental objects that provide object services to client applications in a WebLogic Enterprise domain. You access the environmental objects through a bootstrapping process that accesses the services in a particular WebLogic Enterprise domain.

The following services are provided:

- Object Life Cycle service

The Object Life Cycle service is provided through the FactoryFinder environmental object. The FactoryFinder object is a CORBA object that can be used to locate a factory, which in turn can create object references for CORBA objects. Factories and FactoryFinder objects are implementations of the CORBAservices Life Cycle Service. WebLogic Enterprise applications use the Object Life Cycle service to find object references.

For information about using the Object Life Cycle Service, see “How WebLogic Enterprise CORBA Client and Server Applications Interact” on page 2-16.

- Security service

The Security service is accessed through the SecurityCurrent environmental object. The SecurityCurrent object is used to authenticate a client application into a WebLogic Enterprise domain with the proper security. The WebLogic Enterprise software provides an implementation of the CORBAservices Security Service.

For information about using security, see *Using Security* in the WebLogic Enterprise online documentation.

- Transaction service

The Transaction service is accessed through either the TransactionCurrent environmental object or the UserTransaction object. The TransactionCurrent object allows a client application to participate in a transaction. The WebLogic Enterprise software provides an implementation of the CORBAservices Object Transaction Service (OTS). In addition, the UserTransaction object provides access to the Sun Microsystems, Inc. Java Transaction API (JTA) defined in the `javax.transaction` package.

For information about using transactions, see [Using Transactions](#) in the WebLogic Enterprise online documentation.

- Interface Repository service

The Interface Repository service is accessed through the InterfaceRepository object. The InterfaceRepository object is a CORBA object that contains interface definitions for all the available CORBA interfaces and the factories used to create object references to the CORBA interfaces. The Interface Repository object is used with client applications that use DII.

For information about using DII, see [Creating CORBA Client Applications](#).

The WebLogic Enterprise software provides environmental objects for the following programming environments:

- C++
- Java
- Automation (used by ActiveX client applications)


WebLogic Enterprise Architectural Components

This section provides an introduction to the following architectural components of the WebLogic Enterprise system:

- Bootstrap Object
- IIOP Listener/Handler
- ORB
- TP Framework

Figure 2-3 illustrates the components in a WebLogic Enterprise application.

Figure 2-3 Components in a WebLogic Enterprise Application



Bootstrap Object

The Bootstrap object establishes communication between a client application and a WebLogic Enterprise domain. A domain is simply a way of grouping objects and services together as a management entity. A WebLogic Enterprise domain has at least one IIOP Listener/Handler and is identified by a name. One client application can connect to multiple WebLogic Enterprise domains using different Bootstrap objects.

One of the first things that client applications do after startup is create a Bootstrap object by supplying the host and port of the IIOP Listener/Handler using one of the following URL address formats:


- `//host:port`
- `corbaloc://host:port`
- `corbalocs://host:port`

For more information about the Bootstrap URL address formats, see [Using Security](#) in the WebLogic Enterprise online documentation.

The client application then uses the Bootstrap object to obtain references to the objects in a WebLogic Enterprise domain. Once the Bootstrap object is instantiated, the `resolve_initial_references()` method is invoked by the client application, passing in a `string id`, to obtain a reference to the objects in the domain that provide CORBA services. The valid values for `string id` are `FactoryFinder`, `TransactionCurrent`, `SecurityCurrent`, and `InterfaceRepository`.

Figure 2-4 illustrates how the Bootstrap object works in a WebLogic Enterprise domain.

Figure 2-4 How the Bootstrap Object Works in a WebLogic Enterprise Domain



IIOP Listener/Handler


The IIOP Listener/Handler is a process that receives the client request, which is sent using IIOP, and delivers that request to the appropriate server application. The IIOP Listener/Handler serves as a communication concentrator, providing a critical scalability feature. The IIOP Listener/Handler removes from the server application the burden of maintaining client connections. For information about configuring the IIOP Listener/Handler, see [Administration](#) and the description of the ISL command in the [Commands, Server Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.

ORB

The ORB serves as an intermediary for requests that client applications send to server applications, so that client applications and server applications do not need to contain information about each other. The ORB is responsible for all the mechanisms required to find the implementation that can satisfy the request, to prepare an object's implementation to receive the request, and to communicate the data that makes up the request. The WebLogic Enterprise product provides a C++ ORB and a BEA version of the Java IDL ORB provided with the Java Development Kit (JDK) from Sun Microsystems, Inc.

Figure 2-5 shows the relationship between an ORB, a client application, and a server application.

Figure 2-5 The ORB in a Client/Server Environment



When the client application uses IIOP to send a request to the domain, the ORB performs the following functions:

- Validates each request and its arguments to ensure that the client application supplied all the required arguments.

- Manages the mechanisms required to find the CORBA object that can satisfy the client application's request. To do this, the ORB interacts with the Portable Object Adapter (POA). The POA prepares an object's implementation to receive the request and communicates the data in the request.
- Marshals data. The ORB on the client machine writes the data associated with the request into a standard form. The ORB receives this data and converts it into the format appropriate for the machine on which the server application is running. When the server application sends data back to the client application, the ORB marshals the data back into its standard form and sends it back to the ORB on the client machine.

TP Framework

The TP Framework provides a programming model that achieves high levels of performance while shielding the application programmer from the complexities of the CORBA interfaces. The TP Framework supports the rapid construction of WebLogic Enterprise applications, which makes it easier for application programmers to adhere to design patterns associated with successful TP applications.

The TP Framework interacts with the Portable Object Adapter (POA) and the WebLogic Enterprise application, thus eliminating the need for direct POA calls in an application. In addition, the TP Framework integrates transactions and state management into the WebLogic Enterprise application.

The application programmer uses an application programming interface (API) that automates many of the functions required in a standard CORBA application. The application programmer is responsible only for writing the business logic of the WebLogic Enterprise application and overriding default actions provided by the TP Framework.


The TP Framework API provides routines that perform the following functions required by a CORBA application:

- Initializing the server application and executing startup and shutdown routines
- Creating object references
- Registering and unregistering object factories
- Managing objects and object state

- Tying the server application to WebLogic Enterprise system resources
- Getting and initializing the ORB
- Performing object housekeeping

The TP Framework ensures that the execution of a client request takes place in a coordinated, predictable manner. The TP Framework calls the objects and services available in the WebLogic Enterprise application at the appropriate time, in the correct sequence. In addition, the TP Framework maximizes the reuse of system resources by objects. Figure 2-6 illustrates the TP Framework.

Figure 2-6 The TP Framework



The TP Framework is not a single object, but is rather a collection of objects that work together to manage the CORBA objects that contain and implement your WebLogic Enterprise application's data and business logic.

One of the TP Framework objects is the Server object. The Server object is a user-written programming entity that implements operations that perform tasks such as initializing and releasing the server application; for server applications implemented in C++, the TP Framework instantiates the CORBA objects needed to satisfy a client request.

If a client request that requires an object that is not currently active and in memory in the server application arrives, the TP Framework coordinates all the operations that are required to instantiate the object. This includes coordinating with the ORB and the POA to get the client request to the appropriate object implementation code.

How WebLogic Enterprise CORBA Client and Server Applications Interact


The interaction between WebLogic Enterprise CORBA client and server applications includes the following steps:

1. The server application is initialized.
2. The client application is initialized.
3. The client application authenticates itself to the WebLogic Enterprise domain.
4. The client application obtains a reference to the object needed to execute its business logic.
5. The client application invokes an operation on the CORBA object.

The following topics describe what happens during each step.

Step 1: The Server Application Is Initialized

The system administrator enters the `tmboot` command on a machine in the WebLogic Enterprise domain to start the WebLogic Enterprise server application. The TP Framework invokes the `initialize()` operation in the Server object to initialize the server application.




During the initialization process, the Server object does the following:

1. Gets the Bootstrap object and a reference to the FactoryFinder object.
2. Typically registers any factories with the FactoryFinder object.
3. Optionally gets an object reference to the ORB.
4. Performs any process-wide initialization.

Step 2: The Client Application Is Initialized

During initialization, the client application uses the Bootstrap object in the domain to obtain initial references to the environmental objects available in the domain.



The Bootstrap object returns references to the FactoryFinder, SecurityCurrent, TransactionCurrent, NameService, and InterfaceRepository objects in the WebLogic Enterprise domain.

Step 3: The Client Application Authenticates Itself to the WebLogic Enterprise Domain

If the WebLogic Enterprise domain has a security model in effect, the client application needs to authenticate itself to the WebLogic Enterprise domain before it can invoke any operations in the server application. To authenticate itself to the WebLogic Enterprise domain using Tuxedo authentication, the client application:

1. Uses the Bootstrap object to obtain a reference to the SecurityCurrent object.
2. Invokes the `logon()` operation of the PrincipalAuthenticator object, which is retrieved from the SecurityCurrent object.


Note: For information about using certificate-based authentication, see [Using Security](#) in the WebLogic Enterprise online documentation.

Step 4: The Client Application Obtains a Reference to the Object Needed to Execute Its Business Logic

The client application needs to perform the following steps:

1. Obtain a reference to the factory for the object it needs.


For example, the client application needs a reference to the `SimpleFactory` object. The client application obtains this factory reference from the `FactoryFinder` object, shown in the following figure.




2. Invoke the `SimpleFactory` object to get a reference to the `Simple` object.

If the `SimpleFactory` object is not active, what happens next depends on the programming language in which the server application is implemented:

- In Java, the WebLogic Enterprise system instantiates the `SimpleFactory` object dynamically.
- In C++, the TP Framework instantiates the `SimpleFactory` object by invoking the `Server::create_servant()` method on the `Server` object, shown in the following figure.



3. The TP Framework invokes the `activate_object()` and `find_simple()` operations on the `SimpleFactory` object to get a reference to the `Simple` object, shown in the following figure.




The `SimpleFactory` object then returns the object reference to the `Simple` object to the client application.

Note: Because the TP Framework activates objects by default, the Simpapp sample application does not implicitly use the `activate_object()` operation for the `SimpleFactory` object.


Step 5: The Client Application Invokes an Operation on the CORBA Object

Using the reference to the CORBA object that the factory has returned to the client application, the client application invokes an operation on the object. For example, now that the client application has an object reference to the `Simple` object, the client application can invoke the `to_upper()` operation on it. The instance of the `Simple` object required for the client request is created as shown in the following figure.



If the server application were implemented in Java, the `Simple` object required for the client request is instantiated dynamically by the WebLogic Enterprise system.

The TP Framework invokes the `activate_object()` operation on the `Simple` object and the `SimpleFactory` object to allow the object to initialize any object state necessary, shown in the following figure.



Object state initialization often involves reading durable state information from disk for that object. The TP Framework invokes the operation on the object, returning the response to the client application.

3 The WebLogic Enterprise JavaBeans (EJB) Programming Environment

This topic includes the following sections:

- Overview of the WebLogic Enterprise EJB Programming Environment
- Types of Beans Supported in WebLogic Enterprise
- EJBs and Persistence
- Roles of People Who Develop, Build, Deploy, and Administer EJBs
- Items You Create for an EJB Application
- Tools and Facilities Provided for Building and Deploying EJBs
- EJBs and Failover in the WebLogic Enterprise Environment

Overview of the WebLogic Enterprise EJB Programming Environment

The Enterprise JavaBeans Specification 1.1, published by Sun Microsystems, Inc., defines a component architecture for building distributed, object-oriented business applications in Java. The EJB architecture addresses the development, deployment, and run-time aspects of an enterprise application's lifecycle.

An EJB encapsulates business logic inside a component framework that manages the details of security, transaction, and state management. Low-level details, such as the following, are handled by the EJB container:

- Multithreading
- Resource pooling
- Scaling
- Distributed naming
- Automatic persistence
- Remote invocation
- Transaction boundary management
- Distributed transaction management

This built-in, low-level support allows the EJB to focus on the business problem to be solved.

With the WebLogic Enterprise EJB model, you can write or buy business components (such as invoices, bank accounts, and shipping routes) and, during deployment into a certain project, specify how the component should be used -- which users have access to which methods, whether the container should automatically start a transaction or whether it should inherit the caller's transaction, and so on. In this scenario, an EJB contains the business logic (methods) and the customization needed for a particular application (deployment descriptor), and the EJB will run within any standard implementation of the EJB container. An EJB is, in essence, a distributed object for which transactions and security can be specified declaratively in deployment descriptors.

The spirit of "write once, run anywhere" carries through into EJB: any vendor's EJB container (that conforms to the EJB Specification) can run any third-party EJBs (that also conform to the EJB Specification) to create an application. Nuances of the security mechanisms and specific distributed transaction monitors are entirely abstracted out of the application code (unless the Bean Provider chooses to make such calls explicitly).

Types of Beans Supported in WebLogic Enterprise

With the WebLogic Enterprise system, you can build and deploy standard, portable EJBs. The EJB Specification defines three types of beans as listed and described in Table 3-1.

Table 3-1 Bean Types Specified by the EJB Specification

Bean Type	Description
Stateless session bean	<p>An instance of a stateless session bean has no conversational state for the client that created the instance. This instance is not assigned permanently to the client. The EJB container can maintain a pool of instances and allocate method invocations coming from any client to any available instance (that is, not processing a request for a particular client). Therefore, any instance can receive method invocations from any client, and these requests can be processed on behalf of different transactions and security contexts.</p> <p>The EJB container decides the life of an instance; that is, the container can destroy an instance when resources are required or according to other policies. However, the client decides the life of the reference to the bean. The reference obtained from the bean's home interface is valid until the client destroys it.</p> <p>Note that stateless session beans cannot use the <code>SessionSynchronization</code> interface to synchronize with the starting and stopping of a transaction.</p>

Table 3-1 Bean Types Specified by the EJB Specification (Continued)

Bean Type	Description
Stateful session bean	<p>An instance of a stateful session bean maintains a conversational state for the client that created the instance. Therefore, instances of a stateful session bean are assigned to a particular client and are destroyed only when the client decides to remove the EJB object. Instances of stateful session beans do not survive a crash of the EJB container (which in WebLogic Enterprise spans all the processes in the same group where the bean is deployed) or a redeployment of the bean.</p>
Entity bean	<p>The EJB container can passivate inactive instances to maximize the use of the system resources -- that is, to deactivate the bean with its state saved to be restored at a later time during the bean's reactivation. Stateful session beans can use the <code>SessionSynchronization</code> interface to synchronize with the starting and stopping of a transaction.</p>
	<p>An instance of an entity bean has a unique identity called the primary key. Object references to an entity bean should be usable for a long time and clients should be able to reuse them across server crash or restart. The reference becomes invalid when a client application removes the EJB or when the EJB is reconfigured.</p>
	<p>Note: If a server group crashes, and the System Administrator restarts that group using the same group ID and persistence store, the EJB container can process requests for beans in that group again. The EJB container for stateless session beans spans the entire domain in which the beans are deployed.</p>
	<p>Multiple client applications can access an entity bean instance; the EJB container is responsible for synchronizing the access to the instance.</p>
	<p>Typically, an entity bean has a persistent state, and application designers can choose between managing the persistence directly from the bean (bean-managed) or letting the EJB container manage the persistence (container-managed). In either case, the EJB container determines <i>when</i> an entity bean instance can be passivated (which also triggers the persistent storage of the state of the instance). An entity bean cannot use the <code>SessionSynchronization</code> interface to synchronize with the starting and stopping of a transaction.</p>

EJBs and Persistence

An entity EJB can save its state in any transactional or nontransactional persistent storage, or it can ask the EJB container to save its nontransient instance variables automatically. The WebLogic Enterprise system allows both choices. An EJB that manages its own persistence is referred to as having **bean-managed persistence**; an EJB that delegates to the EJB container the saving and restoring of its state is referred to as having **container-managed persistence**.

You control the persistence characteristics of a bean, such as where its data is maintained in durable storage, in its deployment descriptor; in the case of bean-managed persistence, you implement the specific method invocations in the bean that load and store state.

For more information about development and deployment considerations with regards to persistence, see the following topics:

- Chapter 7, “Developing WebLogic Enterprise EJB Applications.”
- “Development Considerations for EJBs and Persistence” on page 8-11.

Roles of People Who Develop, Build, Deploy, and Administer EJBs

The Enterprise JavaBeans Specification describes the six roles regarding who develops, builds, deploys, and administers an EJB application. These roles are summarized in Table 3-2 to help clarify what needs to be done, by whom, and when during the entire life cycle of an EJB in a way that is consistent with the EJB Specification.

Table 3-2 EJB Roles

Role	Description
Enterprise Bean Provider	<p>The Enterprise Bean Providers (or Bean Providers) produce enterprise beans. Their output is an EJB Java ARchive (EJB JAR) file that contains one or more enterprise beans. The Bean Provider is responsible for:</p> <ul style="list-style-type: none"><li data-bbox="654 453 1231 507">■ The Java classes that implement the enterprise bean's business methods<li data-bbox="654 518 1231 546">■ The definition of the bean's remote and home interfaces<li data-bbox="654 558 1021 585">■ The bean's deployment descriptor <p>The deployment descriptor includes the structural information (for example, the name of the enterprise bean class) of the enterprise bean and declares all the enterprise bean's external dependencies (for example, the names and types of resources that the enterprise bean uses).</p>
Application Assembler	<p>The Application Assembler combines enterprise beans into larger deployable application units. The input to the Application Assembler is one or more EJB JAR files produced by the Bean Providers. The Application Assembler outputs one or more EJB JAR files that contain the enterprise beans along with their application assembly instructions. The Application Assembler has inserted the application assembly instruction into the deployment descriptors.</p> <p>Bean providers cooperate with the Application Assembler to combine EJBs into larger deployable units. In the WebLogic Enterprise environment, creating these larger deployable units is more efficient if the Application Assembler takes into account the scalability and resource management capabilities provided by the WebLogic Enterprise environment. For example, EJBs that access the same resources should be packaged together. The Application Assembler also specifies the security required by the application by associating client role names with the methods of the different beans.</p>

Table 3-2 EJB Roles (Continued)

Role	Description
Deployer	<p>The Deployer uses the EJB container tools to customize one or more EJB JAR files produced by a Bean Provider or Application Assembler so that the beans can run in the corresponding EJB environment. These tools generate the additional classes required to manage the beans. The Deployer is primarily focused on the individual EJBs.</p> <p>In the WebLogic Enterprise environment, the Deployer uses the <code>e jbc</code> command or the WebLogic EJB Deployer for this purpose. These tools can also be used by the Application Assembler to construct an EJB package, which is the EJB JAR file containing all the bean implementations and the assembly instructions. The Deployer also ensures that the security role names defined by the Application Assembler map to existing user groups and accounts that exist in the EJB environment.</p> <p>The Deployer must resolve all the external dependencies declared by the Bean Provider (for example, Deployers must ensure that all resources used by the enterprise beans are present in the operational environment, and they must bind them to the resource manager connection factory references declared in the deployment descriptor), and must follow the application assembly instructions defined by the Application Assembler.</p>
EJB Server Provider	<p>The EJB Server Provider (in the WebLogic Enterprise system, this is BEA) is a specialist in the area of distributed transaction management, distributed objects, and other lower-level, system-level services.</p>

Table 3-2 EJB Roles (Continued)

Role	Description
EJB Container Provider	<p>The EJB Container Provider (in the WebLogic Enterprise system, this is BEA) provides:</p> <ul style="list-style-type: none">■ The deployment tools necessary for the deployment of enterprise beans■ The run-time support for the deployed enterprise beans' instances <p>From the perspective of the enterprise beans, the container is a part of the target operational environment. The container run time provides the deployed enterprise beans with transaction and security management, network distribution of clients, scalable management of resources, and other services that are generally required as part of a manageable server platform.</p>
System Administrator	<p>The System Administrator is responsible for the configuration and administration of the enterprise's computing and networking infrastructure, which includes the EJB server and container. The System Administrator is also responsible for overseeing the well-being of the deployed enterprise bean applications at run time.</p> <p>The System Administrator cooperates with the Deployer to define the environment needed by the application. The System Administrator configures the WebLogic Enterprise domain by defining the different machines, server groups, and other resources needed by the application (for example, JDBC connection pools and XA resource managers).</p> <p>The System Administrator also needs to add any security information needed by the application (for example, new user groups). The administrator is also responsible for monitoring the application and performing any run-time changes needed to adapt the operational environment to failures or other conditions.</p>

Items You Create for an EJB Application

Table 3-3 summarizes all the items you need to create for an EJB application that runs in the WebLogic Enterprise environment, regardless of which role you are assuming, and explains where you can find more information about creating the item.

Table 3-3 Items You Create for an EJB Application

Item	Description	Where to Find More Information
One or more EJBs	The basic beans containing your application's business logic.	"Step 1: Create the EJB" on page 7-8
Deployment descriptor	An XML file, created by one of the following methods, that specifies basic configuration and run-time information relevant to the deployment of the EJBs: <ul style="list-style-type: none">■ DDGenerator command■ WebLogic EJB Deployer■ Manually, using a common text editor	"Step 3: Create the Deployment Descriptor" on page 7-14 and "Step 6: Modify the Deployment Descriptor" on page 7-26
EJB JAR file	A Java ARchive (JAR) file that contains all the Java class files for the EJBs in the application. This file is created initially by the Bean Provider, and is then modified by the Bean Deployer and Application Assembler.	"Step 4: Create a Standard EJB JAR File" on page 7-19 and "Step 7: Package the Components Into a Deployable EJB JAR File" on page 7-28
WebLogic EJB extensions to the deployment descriptor DTD	An XML file, specifying configuration information pertinent to the WebLogic Enterprise environment.	"Step 5: Create the WebLogic EJB Extensions to the Deployment Descriptor DTD" on page 7-21
Module initializer object	A Java object specifying the module initializer class. This entity is optional.	"Step 2: Create the Module Initializer Object" on page 7-13

Tools and Facilities Provided for Building and Deploying EJBs

To help application programmers and deployers build EJBs that fully leverage the WebLogic Enterprise system, the WebLogic Enterprise software provides the tools and facilities listed in Table 3-4.

Table 3-4 Resources for Building and Deploying EJBs


Tool or Facility	Description
<code>ejbc</code> command	Used by application programmers, Application Assemblers, and deployers as a command-line alternative to the WebLogic EJB Deployer to construct a deployable EJB JAR file.
<code>DDGenerator</code> command	Used by the Bean Provider to create the initial deployment descriptor file.
WebLogic EJB Deployer	Used by the Bean Provider, Bean Deployer, and Application Assembler to configure and deploy EJBs for use with your WebLogic Enterprise server. You can use the WebLogic EJB Deployer to: <ul style="list-style-type: none">■ Examine an existing EJB and the configurable properties in its deployment descriptor.■ Modify the properties and save the changes to a file (<code>.xml</code> or <code>.jar</code> format).■ Generate EJB interface classes for a particular WebLogic environment.■ Generate deployment classes for the beans. The WebLogic EJB Deployer is documented in the online help available from that tool's Help menu.
<code>UBBCONFIG</code> file	Used for configuring the EJB container and the Java server in which the EJB container is run and which loads the JVM and other modules needed by the EJB application. (You can also use the TMIB in place of the <code>UBBCONFIG</code> file.)

For more information about deploying and administering EJB applications in the WebLogic Enterprise environment, see Chapter 7, “Developing WebLogic Enterprise EJB Applications.”

EJBs and Failover in the WebLogic Enterprise Environment

The WebLogic Enterprise system provides the following failover characteristics of EJB applications deployed in a WebLogic Enterprise domain. Note that client applications cannot control where EJBs are instantiated. Figure 3-1 shows how, in the instance of a machine crash, failover is managed wholly by the WebLogic Enterprise system.

Figure 3-1 EJB Failover Capabilities in WebLogic Enterprise



Stateless session beans. If the server process hosting a stateless session bean fails, the bean is automatically instantiated in a different server process (on the same server or on another group within the domain), provided that the server process that is capable of supporting the session bean is available.

Entity beans. If one group that hosts one or more entity beans fails and in cases when the client application receives a `RemoteException`, the client application can invoke the `findByPrimaryKey` method to find the home interface for the entity bean, with the specified unique key, on another group in the domain. This works as long as the other group is configured to support that entity bean. Application developers can write client application code within a loop that begins by invoking the `findByPrimaryKey` method; this way, if a group fails, the client application retries the invocation on a different group.

Note that, for bean-managed persistence, the Bean Provider must implement this method explicitly; for container-managed persistence, this method is generated automatically.

Stateful session beans. If one group fails, the administrator must dynamically configure the group on a different machine. For more information, see the [Administration Guide](#) in the WebLogic Enterprise online documentation.

For file-based persistence, recovery depends on whether persistence storage resides on a file system that is still network accessible (for example, an NFS-mounted volume). Because the `persistence-store-directory-root` element in the WebLogic Enterprise EJB extensions to the deployment descriptor DTD specifies the path for persistent storage, the bean's state can be recovered. (Note that this file persistence mechanism is internal to the WebLogic Enterprise system.)

For JDBC-based persistence, the application simply reconnects to the database, as long as the DBMS node is running and the network is accessible to the new node.

Note: In general, you should use JDBC-based persistence for production applications because it is more robust than file-based persistence. File-based persistence is typically appropriate only for development and prototyping purposes.

Part II Developing WebLogic Enterprise CORBA Applications

- Chapter 4. Developing WebLogic Enterprise CORBA Applications
- Chapter 5. Using Security
- Chapter 6. Using Transactions

4 Developing WebLogic Enterprise CORBA Applications

This topic includes the following sections:

- Overview of the Development Process for WebLogic Enterprise CORBA Applications
- The Simpapp Sample Application
- Step 1: Write the OMG IDL Code
- Step 2: Generate Client Stubs and Skeletons
- Step 3: Write the Server Application
- Step 4: Write the Client Application
- Step 5: Create an XA Resource Manager
- Step 6: Create a Configuration File
- Step 7: Create the TUXCONFIG File
- Step 8: Compile the Server Application
- Step 9: Compile the Client Application
- Step 10: Start the WebLogic Enterprise CORBA Application
- Additional WebLogic Enterprise CORBA Sample Applications

For an in-depth discussion of creating WebLogic Enterprise CORBA client and server applications, see the following in the WebLogic Enterprise online documentation:

- [*Creating CORBA Client Applications*](#)
- [*Creating CORBA C++ Server Applications*](#)
- [*Creating CORBA Java Server Applications*](#)

Overview of the Development Process for WebLogic Enterprise CORBA Applications

Table 4-1 outlines the development process for WebLogic Enterprise CORBA applications.


Table 4-1 Development Process for WebLogic Enterprise CORBA Applications

Step	Description
1	Write the Object Management Group (OMG) Interface Definition Language (IDL) code for each CORBA interface you want to use in your WebLogic Enterprise application.
2	Generate the client stubs and the skeletons.
3	Write the server application.
4	Write the client application.
5	Create an XA resource manager.
6	Create a configuration file.
7	Create a TUXCONFIG file.
8	Compile the server application.
9	Compile the client application.
10	Start the WebLogic Enterprise CORBA application.

The steps in the development process are described in the following sections.

Figure 4-1 illustrates the process for developing WebLogic Enterprise CORBA applications.

Figure 4-1 Development Process for WebLogic Enterprise CORBA Applications



* For CORBA Java server applications only

The Simpapp Sample Application


Throughout this topic, the Simpapp sample application is used to demonstrate the development steps. C++ and Java versions of the Simpapp sample application are available.

The server application in the Simpapp sample application provides an implementation of a CORBA object that has the following two methods:

- The `upper()` method accepts a string from the client application and converts the string to uppercase letters.
- The `lower()` method accepts a string from the client application and converts the string to lowercase letters.

Figure 4-2 illustrates how the Simpapp sample application works.

Figure 4-2 Simpapp Sample Application



The source files for the C++ and Java versions of the Simpapp sample application are located in the `\samples\corba\simpapp` and `\samples\corba\simpapp_java` directories of the WebLogic Enterprise software. Instructions for building and running the Simpapp sample applications are in the `Readme.txt` files in the directories. For instructions for building and running the C++ and Java Simpapp sample applications, see [Samples](#) in the WebLogic Enterprise online documentation.

Note: The Simpapp sample applications demonstrate building C++ client and server applications and Java client and server applications. For information about building a simple ActiveX client application, see the [Basic sample application](#) in the WebLogic Enterprise online documentation.

The WebLogic Enterprise product offers a suite of sample applications that demonstrate and aid in the development of WebLogic Enterprise CORBA applications. For an overview of the available sample applications, see [Samples](#) in the WebLogic Enterprise online documentation.

Step 1: Write the OMG IDL Code

The first step in writing a WebLogic Enterprise application is to specify all of the CORBA interfaces and their methods using the Object Management Group (OMG) Interface Definition Language (IDL). An interface definition written in OMG IDL completely defines the CORBA interface and fully specifies each operation's arguments. OMG IDL is a purely declarative language. This means that it contains no implementation details. Operations specified in OMG IDL can be written in and invoked from any language that provides CORBA bindings.

The Simpapp sample application implements the CORBA interfaces listed in Table 4-2.

Table 4-2 CORBA Interfaces for the Simpapp Sample Application

Interface	Description	Operation
SimpleFactory	Creates object references to the Simple object	<code>find_simple()</code>

Table 4-2 CORBA Interfaces for the Simpapp Sample Application (Continued)

Interface	Description	Operation
Simple	Converts the case of a string	to_upper() to_lower()

Listing 4-1 shows the `simple.idl` file that defines the CORBA interfaces in the Simpapp sample application. The same OMG IDL file is used by both the C++ and Java Simpapp sample applications.

Listing 4-1 OMG IDL Code for the Simpapp Sample Application

```
#pragma prefix "beasy.com"

interface Simple
{
    //Convert a string to lower case (return a new string)
    string to_lower(in string val);

    //Convert a string to upper case (in place)
    void to_upper(inout string val);
};

interface SimpleFactory
{
    Simple find_simple();
};
```

Step 2: Generate Client Stubs and Skeletons

The interface specification defined in OMG IDL is used by the IDL compiler to generate client stubs for the client application, and skeletons for the server application. The client stubs are used by the client application for all operation invocations. You use the skeleton, along with the code you write, to create the server application that implements the CORBA objects.

During the development process, use one of the following commands to compile the OMG IDL file and produce client stubs and skeletons for WebLogic Enterprise client and server applications:

- If you are creating C++ client and server applications, use the `idl` command. For a description of the `idl` command, see [Commands, System Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.
- If you are creating Java client and server applications, use the `m3idltojava` command. For a description of the `m3idltojava` command, see [Commands, System Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.

Table 4-3 lists the files that are created by the `idl` command.

Table 4-3 Files Created By the idl Command

File	Default Name	Description
Client stub file	<code>application_c.cpp</code>	Contains generated code for sending a request.
Client stub header file	<code>application_c.h</code>	Contains class definitions for each interface and type specified in the OMG IDL file.
Skeleton file	<code>application_s.cpp</code>	Contains skeletons for each interface specified in the OMG IDL file. During run time, the skeleton maps client requests to the appropriate operation in the server application.
Skeleton header file	<code>application_s.h</code>	Contains the skeleton class definitions.
Implementation file	<code>application_i.cpp</code>	Contains signatures for the methods that implement the operations on the interfaces specified in the OMG IDL file.
Implementation header file	<code>application_i.h</code>	Contains the initial class definitions for each interface specified in the OMG IDL file.

Table 4-4 lists the files that are created by the `m3idltojava` command.

Table 4-4 Files Created By the `m3idltojava` Command

File	Default Name	Description
Base interface class file	<code>interface.java</code>	Contains an implementation of the interface, written in Java. Copy this file to create a new file, and add your business logic to the new file. By convention in the samples and in this document, this file is named <code>interfaceImpl.java</code> . Substitute the actual name of the interface in the filename. This new file is called an <i>object implementation file</i> .
Client stub file	<code>_interfaceStub.java</code>	Contains generated code for sending a request.
Server skeleton file	<code>_interfaceImplBase.java</code>	Contains Java skeletons for each interface specified in the OMG IDL file. During run time, the skeleton maps client requests to the appropriate operation in the Java server application during run time.
Holder class file	<code>interfaceHolder.java</code>	Contains the implementation of the Holder class. The Holder class provides operations for <code>out</code> and <code>inout</code> arguments, which CORBA has, but which do not map exactly to Java.
Helper class file	<code>interfaceHelper.java</code>	Contains the implementation of the Helper class. The Helper class provides auxiliary functionality, notably the <code>narrow</code> method.

Step 3: Write the Server Application

The WebLogic Enterprise software supports C++ and Java server applications. The steps for creating server applications are:

1. Write the methods that implement each interface's operations.
2. Create the server object.

3. Define object activation policies.
4. Create and register a factory.
5. Release the server application.

Writing the Methods That Implement Each Interface's Operations

After you compile the OMG IDL file, you need to write methods that implement the operations for each interface in the file. An implementation file contains the following:

- Method declarations for each operation specified in the OMG IDL file
- Your application's business logic
- Constructors for each interface implementation (implementing these is optional)
- The `activate_object()` and `deactivate_object()` methods (optional)

Within the `activate_object()` and `deactivate_object()` methods, you write code that performs any particular steps related to activating or deactivating the object. For more information, see *Creating CORBA C++ Server Applications* and *Creating CORBA Java Server Applications* in the WebLogic Enterprise online documentation.

You can write the implementation file by hand. However, both the `idl` and `m3idltojava` commands have an option that generates a template for implementation files.

Listing 4-2 includes the C++ implementation of the Simple and SimpleFactory interfaces in the Simpapp sample application.

Listing 4-2 C++ Implementation of the Simple and SimpleFactory Interfaces

```
// Implementation of the Simple_i::to_lower method which converts
// a string to lower case.

char* Simple_i::to_lower(const char* value)
{
    CORBA::String_var var_lower = CORBA::string_dup(value);
```

```
        for (char* ptr = v_lower; ptr && *ptr; ptr++) {
            *ptr = tolower(*ptr);
        }
        return var_lower._retn();
    }

    // Implementation of the Simple_i::to_upper method which converts
    // a string to upper case.

    void Simple_i::to_upper(char*& value)
    {
        CORBA::String_var var_upper = value;
        var_upper = CORBA::string_dup(var_upper.in());
        for (char* ptr = var_upper; ptr && *ptr; ptr++) {
            *ptr = toupper(*ptr);
        }
        value = var_upper._retn();
    }

    // Implementation of the SimpleFactory_i::find_simple method which
    // creates an object reference to a Simple object.

    Simple_ptr SimpleFactory_i::find_simple()
    {
        CORBA::Object_var var_simple_oref =
            TP::create_object_reference(
                _tc_Simple->id(),
                "simple",
                CORBA::NVLList::_nil()
            );
    }
}
```

Listing 4-3 includes the Java implementation of the Simple interface from the Simpapp sample application.

Listing 4-3 Java Implementation of the Simple Interface

```
import com.beasys.Tobj.TP;

/**
 *The SimpleImpl class implements the to_upper and to_lower
 *methods.
 */

public class SimpleImpl extends _SimpleImplBase
{
    /*Converts a string to upper case.*/
}
```

```
public void to_upper(org.omg.CORBA.StringHolder data)
{
    if (data.value == null)
        return;
    data.value = data.value.toUpperCase();
    return;
}
/*Converts a string to lower case.*/
public String to_lower(String data)
{
    if (data == null)
        return null;
    return data.toLowerCase();
}
}
```

Listing 4-4 includes the Java implementation of the SimpleFactory interface from the Simpapp sample application.

Listing 4-4 Java Implementation of the SimpleFactory Interface

```
import com.beasy.Tobj.TP;

/**
 *The SimpleFactoryImpl class provides code to create the Simple
 *object.
 */

public class SimpleFactoryImpl extends _SimpleFactoryImplBase
{
/*Create an object reference to a Simple object*/

    public Simple find_simple()
    {
        org.omg.CORBA.Object simple_oref =
            TP.create_object_reference(
                SimpleHelper.id(), //Repository ID
                "simple",           //object id
                null                //routing criteria
            );
        //Send back the narrowed reference
        return SimpleHelper.narrow(simple_oref);

    };
}
};
```

Creating the Server Object

The Server object performs the following tasks:

- Initializes the server application, including registering factories, allocating resources needed by the server application, and, if necessary, opening an XA resource manager.
- Performs server application shutdown and cleanup procedures.
- In C++ server applications, instantiates CORBA objects needed to satisfy client requests.

In C++ server applications, the Server object is already instantiated and a header file for the Server object is available. You implement methods that initialize and release the server application, and, if desired, create servant objects.

Listing 4-5 includes the C++ code from the Simpapp sample application for the Server object.

Listing 4-5 C++ Server Object

```
static CORBA::Object_var static_var_factory_reference;

// Method to start up the server

CORBA::Boolean Server::initialize(int argc, char* argv[])
{
    // Create the Factory Object Reference

    static_var_factory_reference =
        TP::create_object_reference(
            _tc_SimpleFactory->id(),
            "simple_factory",
            CORBA::NVList::nil()
        );
    // Register the factory reference with the FactoryFinder

    TP::register_factory(
        static_var_factory_reference.in(),
        _tc_SimpleFactory->id()
    );
    return CORBA::TRUE;
}
```

```
}

// Method to shutdown the server

void Server::release()
{
// Unregister the factory.

    try {
        TP::unregister_factory(
            static_var_factory_reference.in(),
            _tc_SimpleFactory->id()
        );
    }
    catch (...) {
        TP::userlog("Couldn't unregister the SimpleFactory");
    }
}
// Method to create servants

Tobj_Servant Server::create_servant(const char*
    interface_repository_id)
{
    if (!strcmp(interface_repository_id,
    _tc_SimpleFactory->id())) {
        return new SimpleFactory_i();
    }
    if (!strcmp(interface_repository_id,
    _tc_Simple->id())) {
        return new Simple_i();
    }
    return 0;
}
```

In Java server applications, you implement the Server object by creating a new class that derives from the `com.beasys.Tobj.Server` class and overrides the `initialize()` and `release()` methods. In the server application code, you can also write a public default constructor for the Server object. When creating Java server applications, you identify the name of the Server object in the Server Description File.

The `create_servant()` method, used in the C++ programming environment of the WebLogic Enterprise product, is not used in the Java environment. In Java, objects are created dynamically, without prior knowledge of the classes being used. In the Java environment of the WebLogic Enterprise product, a servant factory is used to retrieve an implementation class, given the interface repository ID. This information is stored in a server descriptor file. When a method request is received, and no servant is available for the interface, the servant factory looks up the interface and creates an object of the appropriate implementation class.

This collection of the object's implementation and data compose the run-time, active instance of the CORBA object.

When your Java server application starts, the TP Framework creates the Server object specified in the XML file. Then, the TP Framework invokes the `initialize()` method. If the method returns `true`, the server application starts. If the method throws the `com.beasys.TobjS.InitializeFailed` exception, or returns `false`, the server application does not start.

When the server application shuts down, the TP Framework invokes the `release` method on the Server object.

Any command-line options specified in the `CLOPT` parameter for your specific server application in the `SERVERS` section of the WebLogic Enterprise domain's `UBBCONFIG` file are passed to the public `boolean initialize(string[] args)` method as `args`. For more information about passing arguments to the server application, see [Administration Guide](#) in the WebLogic Enterprise online documentation.

Within the `initialize()` method, you can include code that does the following, if applicable:

- Creates and registers factories
- Allocates any machine resources
- Initializes any global variables needed by the server application
- Opens the databases used by the server application
- Opens the XA resource manager

Listing 4-6 includes the Java code from the Simpapp sample application for the Server object.

Listing 4-6 Java Server Object

```
import com.beasys.Tobj.TP;

public class ServerImpl
    extends com.beasys.Tobj.Server
{
    static org.omg.CORBA.Object factory_reference;
```

```
/**Method to start up the server*/

public boolean initialize(String[] args)
{
    try {
        // Create the factory object reference.
        factory_reference = TP.create_object_reference(
            SimpleFactoryHelper.id(),
            "simple_factory",
            null
        );
    }

    // Register the factory reference with the FactoryFinder

        TP.register_factory(
            factory_reference,
            SimpleFactoryHelper.id()
        );
    }

    return true;
}

} catch (Exception e){
    TP.userlog("Couldn't initialize server: " +
    e.getMessage());
    e.printStackTrace();
    return false;
}
}

/** Method to shutdown the server*/

public void release()
{
    try {
        TP.unregister_factory(
            factory_reference,
            SimpleFactoryHelper.id()
        );
    } catch (Exception e){
        TP.userlog("Couldn't unregister the
SimpleFactory: " + e.getMessage());
        e.printStackTrace();
    }
}
}
```

Defining an Object's Activation Policies

As part of server development, you determine what events cause an object to be activated and deactivated by assigning object activation policies, as follows:

- For C++ server applications, specify object activation policies in the Implementation Configuration File (ICF). A template ICF file is created by the `genicf` command.
- For Java server applications, specify object activation policies in the Server Description File, written in Extensible Markup Language (XML).

Note: You also define transaction policies in the ICF and Server Description Files. For information about using transactions in your WebLogic Enterprise CORBA application, see [Using Transactions](#) in the WebLogic Enterprise online documentation.

The WebLogic Enterprise software supports the activation policies listed in Table 4-5.

Table 4-5 Activation Policies

Activation Policy	Description
method	Causes the object to be active only for the duration of the invocation on one of the object's operations. This is the default activation policy.
transaction	Causes the object to be activated when an operation is invoked on it. If the object is activated within the scope of a transaction, the object remains active until the transaction is either committed or rolled back.
process	Causes the object to be activated when an operation is invoked on it, and to be deactivated only when one of the following occurs: <ul style="list-style-type: none">■ The process in which the server application exists is shut down.■ The method <code>TP::deactivateEnable()</code> (C++) or <code>com.beasys.Tobj.TP.deactivateEnable()</code> (Java) has been invoked on the object.

The Simple interface in the Simpapp sample application is assigned the default activation policy of method. For more information about managing object state and defining object activation policies, see [Creating CORBA C++ Server Applications](#) and [Creating CORBA Java Server Applications](#) in the WebLogic Enterprise online documentation.

Creating and Registering a Factory

If your server application manages a factory that you want client applications to be able to locate easily, you need to write the code that registers that factory with the FactoryFinder object.

To write the code that registers a factory managed by your server application, you do the following:

1. Create an object reference to the factory.

You include an invocation to the `create_object_reference()` method, specifying the Interface Repository ID of the factory's OMG IDL interface or the object ID (OID) in string format. In addition, you can specify routing criteria.

2. Register the factory with the WebLogic Enterprise domain.

Use the `register_factory()` method to register the factory with the FactoryFinder object in the WebLogic Enterprise domain. The `register_factory()` method requires the object reference for the factory and a string identifier.

Listing 4-7 includes the code from the C++ Simpapp sample application that creates and registers a factory.

Listing 4-7 C++ Example of Creating and Registering a Factory

```
...
CORBA::Object_var v_reg_oref =
    TP:create_object_reference(
        _tc.SimpleFactory->id(),           //Factory Interface ID
        "simplefactory",                  //Object ID
        CORBA::NVList::_nil());           //Routing Criteria
);
```

```
    TP::register_factory(
        CORBA::Object_var v_reg_oref.in(),
        _tc_SimpleFactory->id(),
    );
    ...
}
```

In Listing 4-7, notice the following:

- `tc.SimpleFactory->id()` specifies the SimpleFactory object's Interface Repository ID by extracting it from its typecode.
- `CORBA::NVList::_nil()` specifies that no routing criteria are used, with the result that an object reference created for the Simple object is routed to the same group as the SimpleFactory object that created the object reference.

Listing 4-8 includes the code from the Java Simpapp sample application that creates and registers a factory.

Listing 4-8 Java Example of Creating and Registering a Factory

```
...
// Save the Simple factory name.
SimpleFName = new String(args[0]);

org.omg.CORBA.Object simple_oref =
    TP.create_object_reference(
        SImpApp.SimpleHelper.id(),    // Repository ID
        SimpleFName,                // Object ID
        null                        // Routing Criteria
    );
// Register the factory reference with the factory finder.

    TP.register_factory(
        fact_oref,                  // factory object reference
        SimpleFName                // factory name
    );
...
}
```

Releasing the Server Application

You need to include code in your server application to perform a graceful shutdown of the server application. The `release()` method is provided for that purpose. Within the `release()` method, you may perform any application-specific cleanup tasks that are specific to the server application, such as:

- Unregistering object factories managed by the server application
- Deallocation resources
- Closing any databases
- Closing an XA resource manager

Once a server application receives a request to shut down, the server application can no longer receive requests from other remote objects. This has implications on the order in which server applications should be shut down, which is an administrative task. For example, do not shut down one server process if a second server process contains an invocation in its `release()` method to the first server process.

During server shutdown, you may want to unregister each of the server application's factories. The invocation of the `unregister_factory()` method should be one of the first actions in the `release()` implementation. The `unregister_factory()` method unregisters the server application's factories. This operation requires the following input arguments:

- The object reference for the factory
- A string identifier, based on the factory object's interface typecode, used to identify the Interface Repository ID of the object's OMG IDL interface

Listing 4-9 includes C++ code that releases a server application and unregistered the factories in the server application.

Listing 4-9 C++ Example of Releasing a WebLogic Enterprise Server Application

```
...
public void release()
{
    TP::unregister_factory(
```

```
        factory_reference.in(),
        SimpleFactoryHelper->id
    );
}

...

```

Listing 4-10 includes Java code that releases a server application and unregistered the factories in the server application.

Listing 4-10 Java Example of Releasing a WebLogic Enterprise Server Application

```
...
/**
 * Method to shutdown the server.
 */
public void release(){}
{
    //This method will only be called if Server.initialize()
    //succeeded, therefore, we know that the factory has been
    //registered with the factory finder.

    //Unregister the factory.
    //Use a try block since cleanup code should not throw
    //exceptions.

    try{
        TP.unregister_factory(
            fact_ref,           //factory object reference
            SimpleFactoryHelper.id() //factory repository id
        );
    }catch (Exception e){
        //Some exception occurred. The call to TP.userlog()
        //will put the message in the ULOG file.
        TP.userlog("Couldn't unregister the SimpleFactory: "
            +e.getMessage());
        e.printStackTrace();
    }
}
...

```

Step 4: Write the Client Application

The WebLogic Enterprise software supports the following types of client applications:

- CORBA C++
- CORBA Java
- CORBA Java applets
- ActiveX

The steps for creating client applications are as follows:

1. Initialize the ORB.
2. Use the Bootstrap environmental object to establish communication with the WebLogic Enterprise domain.
3. Resolve initial references to the FactoryFinder environmental object.
4. Use a factory to get an object reference for the desired CORBA object.
5. Invoke methods on the CORBA object.

Note: For information about creating an ActiveX client application, see *WebLogic Enterprise ActiveX Client Developer's Guide* in the WebLogic Enterprise online documentation.

The client development steps are illustrated in Listing 4-11 and Listing 4-12, which include code from the Simpapp sample application. In the Simpapp sample application, the client application uses a factory to get an object reference to the Simple object and then invokes the `to_upper()` and `to_lower()` methods on the Simple object.

Listing 4-11 C++ Client Application From the Simpapp Sample Application

```
int main(int argc, char* argv[])
{
    try {
        // Initialize the ORB
        CORBA::ORB_var var_orb = CORBA::ORB_init(argc, argv, "");

        // Create the Bootstrap object
        Tobj_Bootstrap bootstrap(var_orb.in(), "");

        // Use the Bootstrap object to find the FactoryFinder
        CORBA::Object_var var_factory_finder_oref =
            bootstrap.resolve_initial_references("FactoryFinder");

        // Narrow the FactoryFinder
        Tobj::FactoryFinder_var var_factory_finder_reference =
            Tobj::FactoryFinder::_narrow
            (var_factory_finder_oref.in());

        // Use the factory finder to find the Simple factory
        CORBA::Object_var var_simple_factory_oref =
            var_factory_finder_reference->find_one_factory_by_id(
                _tc_SimpleFactory->id());
    }

    // Narrow the Simple factory
    SimpleFactory_var var_simple_factory_reference =
        SimpleFactory::_narrow(
            var_simple_factory_reference.in());

    // Find the Simple object
    Simple_var var_simple =
        var_simple_factory_reference->find_simple();

    // Get a string from the user
    cout << "String?";
    char mixed[256];
    cin >> mixed;

    // Convert the string to upper case :
    CORBA::String_var var_upper = CORBA::string_dup(mixed);
    var_simple->to_upper(var_upper.inout());
    cout << var_upper.in() << endl;

    // Convert the string to lower case
    CORBA::String_var var_lower = var_simple->to_lower(mixed);
    cout << var_lower.in() << endl;
}
```

```
        return 0;
    }
}
```

Listing 4-12 Java Client Application From the Simpapp Sample Application

```
public class SimpleClient
{
    public static void main(String args[])
    {
        // Initialize the ORB.
        ORB orb = ORB.init(args, null);

        // Create the Bootstrap object
        Tobj_Bootstrap bootstrap = new Tobj_Bootstrap(orb, "");

        // Use the Bootstrap object to locate the FactoryFinder
        org.omg.CORBA.Object factory_finder_reference =
        bootstrap.resolve_initial_references("FactoryFinder");

        // Narrow the FactoryFinder
        FactoryFinder factory_finder_reference =
        FactoryFinderHelper.narrow(factory_finder_reference);

        // Use the FactoryFinder to find the Simple factory.
        org.omg.CORBA.Object simple_factory_reference =
        factory_finder_reference.find_one_factory_by_id
        (SimpleFactoryHelper.id());

        // Narrow the Simple factory
        SimpleFactory simple_factory_reference =
        SimpleFactoryHelper.narrow(simple_factory_reference);

        // Find the Simple object.
        Simple simple = simple_factory_reference.find_simple();

        // Get a string from the user.
        System.out.println("String?");
        String mixed = in.readLine();

        // Convert the string to upper case.
        org.omg.CORBA.StringHolder buf = new
        org.omg.CORBA.StringHolder(mixed);
        simple.to_upper(buf);
        System.out.println(buf.value);
    }
}
```

```
// Convert the string to lower case.  
String lower = simple.to_lower(mixed);  
System.out.println(lower);  
}  
}
```

Step 5: Create an XA Resource Manager

When using transactions in a WebLogic Enterprise CORBA application, you need to create a server process for the resource manager that interacts with a database on behalf of the WebLogic Enterprise CORBA application. The resource manager you use must conform to the X/OPEN XA specification and you need the following information about the resource manager:

- The name of the structure of type `xa_switch_t` that contains the name of the XA resource manager.
- Flags indicating the capabilities of the XA resource manager and function pointers for the actual XA functions.
- The name of the object files that provide the services of the XA interface.
- The commands needed to open and close the XA resource manager. This information is specified in the `OPENINFO` and `CLOSEINFO` parameters in the `UBBCONFIG` configuration file.

When integrating a new XA resource manager into the WebLogic Enterprise system, the file `$TUXDIR/udataobj/RM` must be updated to include information about the XA resource manager. The information is used to include the correct libraries for the XA resource manager and to automatically and properly set up the interface between the transaction manager and the XA resource manager. The format of this file is as follows:

```
rm_name:rm_structure_name:library_names
```

where `rm_name` is the name of the XA resource manager, `rm_structure_name` is the name of the `xa_switch_t` structure that defines the name of the XA resource manager, and `library_names` is the list of the object files for the XA resource manager. White space (tabs and/or spaces) is allowed before and after each of the values and may be

embedded within the *library_names*. The colon (:) character may not be embedded within any of the values. Lines beginning with a pound sign (#) are treated as comments and are ignored.

Use the `buildtms` command to build a server process for the XA resource manager. The files that result from the `buildtms` command need to be installed in the `$TUXDIR/bin` directory.

For more information about the `buildtms` command, see *Commands, System Processes, and MIB Reference* in the WebLogic Enterprise online documentation.

Step 6: Create a Configuration File

Because the WebLogic Enterprise software offers great flexibility and many options to application designers and programmers, no two applications are alike. An application, for example, may be small and simple (a single client and server running on one machine) or complex enough to handle transactions among thousands of client and server applications. For this reason, for every WebLogic Enterprise CORBA application being managed, the system administrator must provide a configuration file that defines and manages the components (for example, domains, server applications, client applications, and interfaces) of that application.

When system administrators create a configuration file, they are describing the WebLogic Enterprise application using a set of parameters that the WebLogic Enterprise software interprets to create a runnable version of the application. During the setup phase of administration, the system administrator's job is to create a configuration file. The configuration file contains the sections listed in Table 4-6.

Table 4-6 Sections in the Configuration File for WebLogic Enterprise CORBA Applications

Sections in the Configuration File	Description
RESOURCES	Defines defaults (for example, user access and the main administration machine) for the WebLogic Enterprise CORBA application.

Table 4-6 Sections in the Configuration File for WebLogic Enterprise CORBA Applications (Continued)

Sections in the Configuration File	Description
MACHINES	Defines hardware-specific information about each machine running in the WebLogic Enterprise CORBA application.
GROUPS	Defines logical groupings of server applications or CORBA interfaces.
SERVERS	Defines the server application processes (for example, the Transaction Manager) used in the WebLogic Enterprise CORBA application.
SERVICES	Defines parameters for services provided by the WebLogic Enterprise application.
INTERFACES	Defines information about the CORBA interfaces in the WebLogic Enterprise CORBA application.
JDBCONNPOOLS	Describes the pooling of JDBC connections for Java servers.
ROUTING	Defines routing criteria for the WebLogic Enterprise CORBA application.

Listing 4-13 shows the configuration file for the Simpapp sample application.

Listing 4-13 Configuration File for Simpapp Sample Application

```
*RESOURCES
    IPCKEY      55432
    DOMAINID   simpapp
    MASTER      SITE1
    MODEL       SHM
    LDBAL       N

*MACHINES
    "PCWIZ"
    LMID        = SITE1
    APPDIR     = "C:\WLDIR\MY_SIM~1"
    TUXCONFIG  = "C:\WLDIR\MY_SIM~1\results\tuxconfig"
```

```
TUXDIR      = "C:\WLDIR"
MAXWSCLIENTS = 10

*GROUPS
    SYS_GRP
    LMID      = SITE1
    GRPNO    = 1
    APP_GRP
    LMID      = SITE1
    GRPNO    = 2

*Servers
    DEFAULT:
        RESTART = Y
        MAXGEN  = 5
    TMSYSEVT
        SRVGRP  = SYS_GRP
        SRVID   = 1
    TMFFNAME
        SRVGRP  = SYS_GRP
        SRVID   = 2
        CLOPT    = "-A -- -N -M"
    TMFFNAME
        SRVGRP  = SYS_GRP
        SRVID   = 3
        CLOPT    = "-A -- -N"
    TMFFNAME
        SRVGRP  = SYS_GRP
        SRVID   = 4
        CLOPT    = "-A -- -F"
    simple_server
        SRVGRP  = APP_GRP
        SRVID   = 1
        RESTART = N
    ISL
        SRVGRP  = SYS_GRP
        SRVID   = 5
        CLOPT    = "-A -- -n //PCWIZ:2468"

*Services
```

When creating Java server applications, include the `JavaServer` parameter in the `UBBCONFIG` file to start the Java server application. For example:

```
*Servers
.
.
.
JavaServer
```

```
SRVTYPE = JAVA
MODULES = Bankapp.jar
SRVGRP = APP_GRP
SRVID = 2
SYSTEM_ACCESS = FASTPATH
CLOPT = "-A -- -M 10 TellerFactory_1"
RESTART = N
```

If you are using an XA resource manager, use the `JavaServerXA` parameter in place of the `JavaServer` parameter to associate the XA resource manager with a specified server group. You need to include the information to open and close the resource manager in the `OPENINFO` and `CLOSEINFO` parameters in the `GROUPS` section of the `UBBCONFIG` file. The information needed to open and close the resource manager should be provided by the manufacturer of the resource manager.

Step 7: Create the TUXCONFIG File

There are two forms of the configuration file:

- An ASCII version of the file, created and modified with any editor. Throughout the WebLogic Enterprise documentation, the ASCII version of the configuration file is referred to as the `UBBCONFIG` file. The configuration file may, in fact, be given any filename.
- The `TUXCONFIG` file, a binary version of the `UBBCONFIG` file created using the `tmloadcf` command. When the `tmloadcf` command is executed, the environment variable `TUXCONFIG` must be set to the name and directory location of the `TUXCONFIG` file. The `tmloadcf` command converts the configuration file to binary form and writes it to the location specified in the command.

For more information about the `tmloadcf` command, see [*Commands, System Processes, and MIB Reference*](#) in the WebLogic Enterprise online documentation.

Step 8: Compile the Server Application

You use the `buildobjserver` command to compile and link C++ server applications. The `buildobjserver` command has the following format:

```
buildobjserver [-o servername] [options]
```

In the `buildobjserver` command syntax:

- `-o servername` represents the name of the server application to be generated by this command.
- `options` represents the command-line options to the `buildobjserver` command.

When creating Java server applications, use the `javac` compiler to create the bytecodes for all the class files that comprise your WebLogic Enterprise CORBA application. This set of files includes the `*.java` source files generated by the `m3idl to java` compiler, plus the object implementation files and server class files you created.

You use the `buildjavaserver` command to build a Java ARchive (JAR) file and link the Java server applications. The `buildjavaserver` command has the following format:

```
buildjavaserver [-s searchpath] input_file.xml
```

In the `buildjavaserver` command syntax:

- `-s searchpath` is used to locate the classes and packages when building the archive. If this optional value is not specified, it defaults to the value of the `CLASSPATH` environment variable.
- `input_file` is the name of the XML Server Description File.

You then need to specify in the `APPPDIR` system environment variable the location of the JAR file for your Java server application. On Windows NT systems, this directory must be on a local drive (not a networked drive). On Solaris systems, the directory can be local or remote.

Step 9: Compile the Client Application

The final step in the development of the CORBA client application is to produce the executable client application. To do this, you need to compile the code and then link against the client stub.

When creating CORBA C++ client applications, use the `buildobjclient` command to construct a WebLogic Enterprise client application executable. The command combines the client stubs for interfaces that use static invocation, and the associated header files, with the standard WebLogic Enterprise libraries to form a client executable. For the syntax of the `buildobjclient` command, see [Commands, System Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.

When creating CORBA Java client applications, see your Java ORB's documentation for information about building client executables. You need to include the `wledir\udataobj\java\jdk\m3envobj.jar` file in your `CLASSPATH` when you compile the CORBA Java client application. The `m3envobj.jar` file contains the Java classes for the WebLogic Enterprise environmental objects.

Step 10: Start the WebLogic Enterprise CORBA Application

Use the `tmboot` command to start the server processes in your WebLogic Enterprise CORBA application. The WebLogic Enterprise CORBA application is usually booted from the machine designated as the `MASTER` in the `RESOURCES` section of the `UBBCONFIG` file.

For the `tmboot` command to find executables, the WebLogic Enterprise system processes must be located in `$TUXDIR/bin`. Server applications should be in `APPDIR`, as specified in the configuration file.

When booting server applications, the `tmboot` command uses the `CLOPT`, `SEQUENCE`, `SRVGRP`, `SRVID`, and `MIN` parameters from the configuration file. Server applications are booted in the order in which they appear in the configuration file.

For more information about using the `tmboot` command, see [Commands, System Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.

Additional WebLogic Enterprise CORBA Sample Applications

Sample applications demonstrate the tasks involved in developing a WebLogic Enterprise CORBA application, and provide sample code that can be used by client and server programmers to build their own WebLogic Enterprise CORBA application. Code from the sample applications are used throughout the information topics in the WebLogic Enterprise product to illustrate the development and administrative steps. For information about building and running the sample applications, see [Samples](#) in the WebLogic Enterprise online documentation.

Table 4-7 describes the additional WebLogic Enterprise CORBA sample applications.

Table 4-7 The WebLogic Enterprise CORBA Sample Applications

WebLogic Enterprise CORBA Sample Application	Description
Simpapp	Provides a C++ client application and a C++ server application. The C++ server application contains two operations that manipulate strings received from the C++ client application.
Java Simpapp	Provides a Java client application and a Java server application. The Java server application contains two operations that manipulate strings received from the Java client application.

Table 4-7 The WebLogic Enterprise CORBA Sample Applications (Continued)

WebLogic Enterprise CORBA Sample Application	Description
Basic	Describes how to develop WebLogic client and server applications and configure the WebLogic Enterprise application. Building C++ server applications and CORBA C++, CORBA Java, and ActiveX client applications are demonstrated.
Security	Demonstrates adding Tuxedo authentication to a WebLogic Enterprise application. For information about building and running the Security sample application, see Using Security in the WebLogic Enterprise online documentation.
Transactions	Adds transactional objects to the C++ server application and client applications in the Basic sample application. The Transactions sample application demonstrates how to use the Implementation Configuration File (ICF) to define transaction policies for CORBA objects. For information about building and running the Transactions sample application, see Using Transactions in the WebLogic Enterprise online documentation.
Wrapper	Demonstrates how to wrap an existing BEA Tuxedo application as a CORBA object.
Production	Demonstrates replicating server applications, creating stateless objects, and implementing factory-based routing in server applications.
JDBC Bankapp	Implements an automatic teller machine (ATM) interface and uses Java Database Connectivity (JDBC) to access a database that stores account and customer information. For information about building and running the JDBC Bankapp sample application, see Using Transactions in the WebLogic Enterprise online documentation.

Table 4-7 The WebLogic Enterprise CORBA Sample Applications (Continued)

WebLogic Enterprise CORBA Sample Application	Description
XA Bankapp	Implements the same ATM interface as JDBC Bankapp; however, XA Bankapp uses a database XA library to demonstrate using the Transaction Manager to coordinate transactions. For information about building and running the XA Bankapp sample application, see <i>Using Transactions</i> in the WebLogic Enterprise online documentation.
Secure Simpapp	Implements the necessary development and administrative changes to the Simpapp sample application to support certificate-based authentication. Java and C++ versions are provided. For information about building and running the Secure Simpapp sample application, see <i>Using Security</i> in the WebLogic Enterprise online documentation.
Introductory Events	Demonstrates how to use joint client/server applications and callback objects to implement events in a WebLogic Enterprise CORBA application. The C++ version uses the BEA Simple Events API and the Java version uses the CosNotification API. For information about building and running the Introductory Events sample application, see <i>Using the Notification Service</i> in the WebLogic Enterprise online documentation.
Advanced Events	Provides a more complex implementation of events in a WebLogic Enterprise CORBA application with transient and persistent subscriptions and data filtering. The C++ version uses the BEA Simple Events API and the Java version uses the CosNotification API. For information about building and running the Advanced Events sample application, see <i>Using the Notification Service</i> in the WebLogic Enterprise online documentation.

5 Using Security

This topic includes the following sections:

- Overview of the Security Service
- How Security Works
- The Security Sample Application
- Development Steps

Note: This chapter describes using username/password authentication. For a complete description of all the security features available in the WebLogic Enterprise product and instructions for implementing the security features, see [*Using Security*](#) in the WebLogic Enterprise online documentation.

Overview of the Security Service

The WebLogic Enterprise product offers a security model based on the CORBAAservices Security Service. The WebLogic Enterprise security model implements the authentication portion of the CORBAAservices Security Service.

Security information is defined on a domain basis. The security level for the domain is defined in the configuration file. Client applications use the SecurityCurrent object to provide the necessary authentication information to log on to the WebLogic Enterprise domain.

The following levels of authentication are provided:

- **TOBJ_NOAUTH**

No authentication is needed; however, the client application may still authenticate itself, and may specify a username and a client application name, but no password.

- **TOBJ_SYSAUTH**

The client application must authenticate itself to the WebLogic Enterprise domain and must specify a username, client application name, and application password.

- **TOBJ_APPAUTH**

In addition to the TOBJ_SYSAUTH information, the client application must provide application-specific information. If the default WebLogic Enterprise authentication service is used in the application configuration, the client application must provide a user password; otherwise, the client application provides authentication data that is interpreted by the custom authentication service in the application.


Note: If a client application is not authenticated and the security level is TOBJ_NOAUTH, the IIOP Listener/Handler of the WebLogic Enterprise domain registers the client application with the username and client application name sent to the IIOP Listener/Handler.

In the WebLogic Enterprise software, only the PrincipalAuthenticator and Credentials properties on the SecurityCurrent object are supported. For a description of the SecurityLevel1::Current and SecurityLevel2::Current interfaces, see the C++ and Java topics in *Commands, System Processes, and MIB Reference* in the WebLogic Enterprise online documentation.

How Security Works

Figure 5-1 illustrates how security works in a WebLogic Enterprise domain.

Figure 5-1 How Security Works in a WebLogic Enterprise Domain



The steps are as follows:

1. The client application uses the Bootstrap object to return an object reference to the SecurityCurrent object for the WebLogic Enterprise domain.
2. The client application obtains the PrincipalAuthenticator.
3. The client application uses the `Tobj::PrincipalAuthenticator::get_auth_type()` method to get the authentication level for the WebLogic Enterprise domain.
4. The proper authentication level is returned to the client application.
5. The client application uses the `Tobj::PrincipalAuthenticator::logon()` method to log on to the WebLogic Enterprise domain with the proper authentication information.


The Security Sample Application

The Security sample application demonstrates username/password authentication. The Security sample application requires each student using the application to have an ID and a password. The Security sample application works in the following manner:

- The client application has a `logon()` operation. This operation invokes operations on the `PrincipalAuthenticator` object, which is obtained as part of the process of logging on to access the domain.
- The server application implements a `get_student_details()` operation on the `Registrar` object to return information about a student. After the user is authenticated, `logon` is complete and the `get_student_details()` operation accesses the student information in the database to obtain the student information needed by the client `logon` operation.
- The database in the Security sample application contains course and student information.

Figure 5-2 illustrates the Security sample application.

Figure 5-2 Security Sample Application



The source files for the Security sample application are located in the `\samples\corba\university` directory in the WebLogic Enterprise software. For information about building and running the Security sample application, see [Using Security](#) in the WebLogic Enterprise online documentation.

Development Steps

Table 5-1 lists the development steps for writing a WebLogic Enterprise CORBA application that has username/password authentication security.

Table 5-1 Development Steps for WebLogic Enterprise CORBA Applications That Have Security

Step	Description
1	Define the security level in the configuration file.
2	Write the CORBA client application.

Step 1: Define the Security Level in the Configuration File

The security level for a WebLogic Enterprise domain is defined by setting the SECURITY parameter in the RESOURCES section of the configuration file to the desired security level. Table 5-2 lists the options for the SECURITY parameter.

Table 5-2 Options for the SECURITY Parameter

Option	Definition
NONE	No security is implemented in the domain. This option is the default. This option maps to the TOBJ_NOAUTH level of authentication.
APP_PW	Requires that client applications provide an application password during initialization. The <code>tmloadcf</code> command prompts for an application password. This option maps to the TOBJ_APPAUTH level of authentication.
USER_AUTH	Requires an application password and performs a per-user authentication during the initialization of the client application. This option maps to the TOBJ_SYSAUTH level of authentication.

In the Security sample application, the SECURITY parameter is set to APP_PW for application-level security. For information about adding security to a WebLogic Enterprise CORBA application, see [Using Security](#) in the WebLogic Enterprise online documentation.

Step 2: Write the CORBA Client Application

Write client application code that does the following:

1. Uses the Bootstrap object to obtain a reference to the SecurityCurrent object for the specific WebLogic Enterprise domain.
2. Gets the PrincipalAuthenticator object from the SecurityCurrent object.
3. Uses the `get_auth_type()` operation of the PrincipalAuthenticator object to return the type of authentication expected by the WebLogic Enterprise domain.

Listing 5-1 and Listing 5-2 include the portions of the CORBA C++ and CORBA Java client applications in the Security sample application that illustrate the development steps for security.

Listing 5-1 Example of Security in a CORBA C++ Client Application

```
CORBA::Object_var var_security_current_oref =
    bootstrap.resolve_initial_references("SecurityCurrent");
SecurityLevel2::Current_var var_security_current_ref =
    SecurityLevel2::Current::_narrow(var_security_current_oref.in());

//Get the PrincipalAuthenticator
SecurityLevel2::PrincipalAuthenticator_var var_principal_authenticator_oref =
    var_security_current_ref->principal_authenticator();
//Narrow the PrincipalAuthenticator
Tobj::PrincipalAuthenticator_var var_bea_principal_authenticator =
    Tobj::PrincipalAuthenticator::_narrow(
        var_principal_authenticator_oref.in());

//Determine the security level
Tobj::AuthType auth_type = var_bea_principal_authenticator->get_auth_type();
Security::AuthenticationStatus status = var_bea_principalauthenticator->logon(
    user_name,
    client_name,
```

```
system_password,  
user_password,  
0);
```

Listing 5-2 Example of Security in a CORBA Java Client Application

```
org.omg.CORBA.Object SecurityCurrentObj =  
    gBootstrapObjRef.resolve_initial_references("SecurityCurrent");  
org.omg.SecurityLevel2.Current secCur =  
    org.omg.SecurityLevel2.CurrentHelper.narrow(secCurObj);  
  
//Get the PrincipalAuthenticator  
org.omg.SecurityLevel2.PrincipalAuthenticator authlevel2 =  
    secCur.principal_authenticator();  
//Narrow the PrincipalAuthenticator  
com.beasys.Tobj.PrincipalAuthenticatorObjRef gPrinAuthObjRef =  
    (com.beasys.Tobj.PrincipalAuthenticator)  
    org.omg.SecurityLevel2.PrincipalAuthenticatorHelper.narrow(authlevel2);  
  
//Determine the security level  
com.beasys.Tobj.AuthType authType = gPrinAuthObjRef.get_auth_type();  
  
org.omg.Security.AuthenticationStatus status = gPrinAuthObjRef.logon  
    (gUserName, ClientName, gSystemPassword, gUserPassword, 0);
```

6 Using Transactions

This topic includes the following sections:

- Overview of the Transaction Service
- What Happens During a Transaction
- Transactions Sample Application
- Development Steps

Note: This topic describes using the C++ interface to the CORBAservices Object Transaction service. For a complete description of all the transaction features available in the WebLogic Enterprise product and instructions for implementing the transaction features, see [Using Transactions](#) in the WebLogic Enterprise online documentation.

Overview of the Transaction Service

One of the most fundamental features of the WebLogic Enterprise product is transaction management. Transactions are a means to guarantee that database transactions are completed accurately and that they take on all the **ACID properties** (atomicity, consistency, isolation, and durability) of a high-performance transaction. The WebLogic Enterprise system protects the integrity of your transactions by providing a complete infrastructure for ensuring that database updates are done accurately, even across a variety of resource managers.

The WebLogic Enterprise system includes the following:

- The CORBAservices Object Transaction Service (OTS) and the Java Transaction Service (JTS)

The WebLogic Enterprise product provides a C++ interface to the OTS and a Java interface to the OTS via the JTS. The JTS is the Sun Microsystems, Inc. Java interface for transaction services, and is based on the OTS. The OTS and the JTS are accessed through the TransactionCurrent environmental object. For information about using the TransactionCurrent environmental object, see [*CORBA C++ Programming Reference*](#) or [*CORBA Java Programming Reference*](#) in the WebLogic Enterprise online documentation.

- The Sun Microsystems, Inc. Java Transaction API (JTA)

Only the application-level demarcation interface (`javax.transaction.UserTransaction`) is supported. For information about JTA, refer to the following:

- The `javax.transaction` package description in the [*API Javadoc*](#).
- The Java Transaction API specification, published by Sun Microsystems, Inc. and available from the Sun Microsystems, Inc. Web site.

OTS, JTS, and JTA each provide the following support for your business transactions:


- Creates a global transaction identifier when a client application initiates a transaction.
- Works with the TP Framework to track objects that are involved in a transaction and, therefore, need to be coordinated when the transaction is ready to commit.
- Notifies the resource managers—which are, most often, databases—when they are accessed on behalf of a transaction. Resource managers then lock the accessed records until the end of the transaction.
- Orchestrates the two-phase commit when the transaction completes, which ensures that all the participants in the transaction commit their updates simultaneously. It coordinates the commit with any databases that are being updated using The Open Group XA protocol. Almost all relational databases support this standard.
- Executes the rollback procedure when the transaction must be stopped.

- Executes a recovery procedure when failures occur. It determines which transactions were active in the machine at the time of the crash, and then determines whether the transaction should be rolled back or committed.

What Happens During a Transaction

Figure 6-1 illustrates how transactions work in a WebLogic Enterprise CORBA application.

Figure 6-1 How Transactions Work in a WebLogic Enterprise CORBA Application



A basic transaction works in the following way:

1. The client application uses the `Bootstrap` object to return an object reference to the `TransactionCurrent` object for the WebLogic Enterprise domain.
2. A client application begins a transaction using the `Tobj::TransactionCurrent::begin()` method, and issues a request to the CORBA interface through the TP Framework. All operations on the CORBA interface execute within the scope of a transaction.
 - If a call to any of these operations raises an exception (either explicitly or as a result of a communication failure), the exception can be caught and the transaction can be rolled back.
 - If no exceptions occur, the client application commits the current transaction using the `Tobj::TransactionCurrent::commit()` method. This method ends the transaction and starts the processing of the operation. The transaction is committed only if all of the participants in the transaction agree to commit.
3. The `Tobj::TransactionCurrent::commit()` method causes the TP Framework to call the Transaction Manager to complete the transaction.
4. The Transaction Manager updates the database.

Transactions Sample Application

In the Transactions sample application, the operation of registering for courses is executed within the scope of a transaction. The transaction model used in the Transactions sample application is a combination of the conversational model and the model in which a single client invocation invokes multiple individual operations on a database.

The Transactions sample application works in the following way:

1. Students submit a list of courses for which they want to be registered.
2. For each course in the list, the server application checks whether:
 - The course is in the database


- The student is already registered for a course
- The student exceeds the maximum number of credits the student can take

3. One of the following occurs:

- If the course meets all the criteria, the server application registers the student for the course.
- If the course is not in the database or if the student is already registered for the course, the server application adds the course to a list of courses for which the student could not be registered. After processing all the registration requests, the server application returns the list of courses for which registration failed. The client application can then choose to either commit the transaction (thereby registering the student for the courses for which registration request succeeded) or to roll back the transaction (thus, not registering the student for any of the courses).
- If the student exceeds the maximum number of credits the student can take, the server application returns a `TooManyCredits` user exception to the client application. The client application provides a brief message explaining that the request was rejected. The client application then rolls back the transaction.

Figure 6-2 illustrates how the Transactions sample application works.

Figure 6-2 Transactions Sample Application



Part of a Transaction

The Transactions sample application shows two ways in which a transaction can be rolled back:

- **Nonfatal.** If the registration for a course fails because the course is not in the database, or because the student is already registered for the course, the server application returns the numbers of those courses to the client application. The decision to roll back the transaction lies with the user of the client application.
- **Fatal.** If the registration for a course fails because the student exceeds the maximum number of credits he or she can take, the server application generates a CORBA exception and returns it to the client application. The decision to roll back the transaction also lies with the client application.

Note: For information about how transactions are implemented in WebLogic Enterprise CORBA Java applications, see the description of the XA Bankapp sample application in [Using Transactions](#) in the WebLogic Enterprise online documentation.

Development Steps

This topic describes the development steps for writing a WebLogic Enterprise CORBA application that includes transactions. Table 6-1 lists the development steps.

Table 6-1 Development Steps for WebLogic Enterprise CORBA Applications That Have Transactions

Step	Description
1	Write the OMG IDL code for the transactional CORBA interface.
2	Define the transaction policies for the CORBA interface in the Implementation Configuration file (ICF) for C++ WebLogic Enterprise CORBA applications, or in the Server Description File for Java WebLogic Enterprise CORBA client applications.
3	Write the client application.
4	Write the server application.
5	Create a configuration file.

The Transactions sample application is used to demonstrate these development steps. The source files for the Transactions sample application are located in the `\samples\corba\university` directory of the WebLogic Enterprise software. For information about building and running the Transactions sample application, see [Samples](#) in the WebLogic Enterprise online documentation.

The XA Bankapp sample application demonstrates how to use transactions in Java WebLogic Enterprise CORBA applications. The source files for the XA Bankapp sample application are located in the `\samples\corba\bankapp_java` directory of the WebLogic Enterprise software. For information about building and running the XA Bankapp sample application, see [Samples](#) in the WebLogic Enterprise online documentation.

Step 1: Write the OMG IDL Code

You need to specify interfaces involved in transactions in Object Management Group (OMG) Interface Definition Language (IDL) just as you would any other CORBA interface. You must also specify any user exceptions that may occur from using the interface.

For the Transactions sample application, you would define in OMG IDL the `Registrar` interface and the `register_for_courses()` operation. The `register_for_courses()` operation has a parameter, `NotRegisteredList`, which returns to the client application the list of courses for which registration failed. If the value of `NotRegisteredList` is empty, the client application commits the transaction. You also need to define the `TooManyCredits` user exception.

Listing 6-1 includes the OMG IDL code for the Transactions sample application.

Listing 6-1 OMG IDL Code for the Transactions Sample Application

```
#pragma prefix "beasy.com"
module University

{
    typedef unsigned long CourseNumber;
    typedef sequence<CourseNumber> CourseNumberList;

    struct CourseSynopsis
    {
        CourseNumber    course_number;
        string          title;
    };
    typedef sequence<CourseSynopsis> CourseSynopsisList;

    interface CourseSynopsisEnumerator
    {
        //Returns a list of length 0 if there are no more entries
        CourseSynopsisList get_next_n(
            in unsigned long number_to_get, // 0 = return all
            out unsigned long number_remaining
        );
        void destroy();
    };
}
```

```
typedef unsigned short Days;
const Days MONDAY      = 1;
const Days TUESDAY     = 2;
const Days WEDNESDAY   = 4;
const Days THURSDAY    = 8;
const Days FRIDAY      = 16;

//Classes restricted to same time block on all scheduled days,
//starting on the hour

struct ClassSchedule
{
    Days          class_days; // bitmask of days
    unsigned short start_hour; // whole hours in military time
    unsigned short duration;  // minutes
};

struct CourseDetails
{
    CourseNumber course_number;
    double        cost;
    unsigned short number_of_credits;
    ClassSchedule class_schedule;
    unsigned short number_of_seats;
    string        title;
    string        professor;
    string        description;
};
typedef sequence<CourseDetails> CourseDetailsList;
typedef unsigned long StudentId;

struct StudentDetails
{
    StudentId      student_id;
    string         name;
    CourseDetailsList registered_courses;
};

enum NotRegisteredReason
{
    AlreadyRegistered,
    NoSuchCourse
};

struct NotRegistered
{
    CourseNumber    course_number;
    NotRegisteredReason not_registered_reason;
};

typedef sequence<NotRegistered> NotRegisteredList;
```

6 Using Transactions

```
exception TooManyCredits
{
    unsigned short maximum_credits;
};

//The Registrar interface is the main interface that allows
//students to access the database.
interface Registrar
{
    CourseSynopsisList
    get_coursesSynopsis(
        in string                               search_criteria,
        in unsigned long                         number_to_get,
        out unsigned long                        number_remaining,
        out CourseSynopsisEnumerator rest
    );

    CourseDetailsList get_coursesDetails(in CourseNumberList
    courses);
    StudentDetails get_student_details(in StudentId student);
    NotRegisteredList register_for_courses(
        in StudentId      student,
        in CourseNumberList courses
    ) raises (
        TooManyCredits
    );
};

// The RegistrarFactory interface finds Registrar interfaces.

interface RegistrarFactory
{
    Registrar find_registrar(
    );
};


```

Step 2: Define Transaction Policies for the Interfaces

Transaction policies are used on a per-interface basis. During design, it is decided which interfaces within a WebLogic Enterprise application will handle transactions. The transaction policies are:

Transaction Policy	Description
always	The interface must always be part of a transaction. If the interface is not part of a transaction, a transaction will be automatically started by the TP Framework.
ignore	The interface is not transactional; however, requests made to this interface within a scope of a transaction are allowed. The AUTOTRAN parameter, specified in the UBBCONFIG file for this interface, is ignored.
never	The interface is not transactional. Objects created for this interface can never be involved in a transaction. The WebLogic Enterprise system generates an exception (INVALID_TRANSACTION) if an interface with this policy is involved in a transaction.
optional	The interface may be transactional. Objects can be involved in a transaction if the request is transactional. This transaction policy is the default.

During development, you decide which interfaces will execute in a transaction by assigning transaction policies, as follows:

- For C++ server applications, you specify transaction policies in the Implementation Configuration File (ICF). A template ICF file is created by the `genicf` command.
- For Java server applications, you specify transaction policies in the Server Description File, written in Extensible Markup Language (XML).

In the Transactions sample application, the transaction policy of the `Registrar` interface is set to `always`.

Step 3: Write the CORBA Client Application

The CORBA client application needs code that performs the following tasks:

1. Obtains a reference to the `TransactionCurrent` object from the `Bootstrap` object.

2. Begins a transaction by invoking the `Tobj::TransactionCurrent::begin()` operation on the `TransactionCurrent` object.
3. Invokes operations on the object. In the `Transactions` sample application, the client application invokes the `register_for_courses()` operation on the `Registrar` object, passing a list of courses.

Listing 6-2 illustrates the portion of the CORBA C++ client applications in the `Transactions` sample application that illustrates the development steps for transactions.

For an example of a CORBA Java client application that uses transactions, see the XA Bankapp sample application in [Guide to the Java Sample Applications](#) in the WebLogic Enterprise online documentation.

Listing 6-2 Transactions Code for CORBA C++ Client Applications

```
CORBA::Object_var var_transaction_current_oref =
    Bootstrap.resolve_initial_references("TransactionCurrent");
CosTransactions::Current_var transaction_current_oref=
    CosTransactions::Current::_narrow(var_transaction_current_oref.in());
//Begin the transaction
var_transaction_current_oref->begin();
try {
    //Perform the operation inside the transaction
    pointer_Registar_ref->register_for_courses(student_id, course_number_list);
    ...
    //If operation executes with no errors, commit the transaction:
    CORBA::Boolean report_heuristics = CORBA::TRUE;
    var_transaction_current_ref->commit(report_heuristics);
}
catch (...) {
    //If the operation has problems executing, rollback the
    //transaction. Then throw the original exception again.
    //If the rollback fails, ignore the exception and throw the
    //original exception again.
    try {
        var_transaction_current_ref->rollback();
    }
    catch (...) {
        TP::userlog("rollback failed");
    }
    throw;
}
```

Step 4: Write the Server Application

When using transactions in server applications, you need to write methods that implement the interface's operations. In the Transactions sample application, you would write a method implementation for the `register_for_courses()` operation.

If your WebLogic Enterprise CORBA application uses a database, you need to include code in the server application that opens and closes an XA resource manager. These operations are included in the `Server::initialize()` and `Server::release()` operations of the Server object.

Listing 6-3 shows the portion of the code for the Server object in the Transactions sample application that opens and closes the XA resource manager.

Note: For a complete example of a C++ server application that implements transactions, see the Transactions sample application in [Using Transactions](#) in the WebLogic Enterprise online documentation.

For an example of a Java server application that implements transactions, see the description of the XA Bankapp sample application in [Using Transactions](#) in the WebLogic Enterprise online documentation.

Listing 6-3 C++ Server Object in Transactions Sample Application

```
CORBA::Boolean Server::initialize(int argc, char* argv[])
{
    TRACE_METHOD("Server::initialize");
    try {
        open_database();
        begin_transactional();
        register_fact();
        return CORBA_TRUE;
    }
    catch (CORBA::Exception& e) {
        LOG("CORBA exception : " << e);
    }
    catch (SamplesDBException& e) {
        LOG("Can't connect to database");
    }
    catch (...) {
        LOG("Unexpected exception");
    }
    cleanup();
}
```

6 Using Transactions

```
        return CORBA_FALSE;
    }

void Server::release()
{
    TRACE_METHOD("Server::release");
    cleanup();
}

static void cleanup()
{
    unregister_factory();
    end_transactional();
    close_database();
}
//Utilities to manage transaction resource manager

CORBA::Boolean s_became_transactional = CORBA_FALSE;
static void begin_transactional()
{
    TP::open_xa_rm();
    s_became_transactional = CORBA_TRUE;
}
static void end_transactional()
{
    if(!s_became_transactional){
        return//cleanup not necessary
    }
    try {
        TP::close_xa_rm ();
    }
    catch (CORBA::Exception& e) {
        LOG("CORBA Exception : " << e);
    }
    catch (...) {
        LOG("unexpected exception");
    }
    s_became_transactional = CORBA_FALSE;
}
```

Step 5: Create a Configuration File

You need to add the following information to the configuration file for a transactional WebLogic Enterprise CORBA application.

- In the SERVERS section:
 - Define a server group that includes both the server application that includes the interface and the server application that manages the database. This server group needs to be specified as transactional.
 - Replace JavaServer with JavaServerXA to associate the XA resource manager with a specified server group. (JavaServer uses the null RM.)
- In the OPENINFO and CLOSEINFO parameters of the GROUPS section, include information to open and close the XA resource manager for the database. You obtain this information from the product documentation for your database. Note that the default version of the com.beasys.Tobj.Server.initialize() operation automatically opens the resource manager.
- Include the pathname to the transaction log (TLOG) in the TLOGDEVICE parameter. For more information about the transaction log, see [Administration](#) in the WebLogic Enterprise online documentation.

Listing 6-4 includes the portions of the configuration file that define this information for the Transactions sample application.

Listing 6-4 Configuration File for Transactions Sample Application

```
*RESOURCES
    IPCKEY      55432
    DOMAINID   university
    MASTER      SITE1
    MODEL       SHM
    LDBAL       N
    SECURITY    APP_PW

*MACHINES
    BLOTTO
    LMID = SITE1
    APPDIR = C:\TRANSACTION_SAMPLE
    TUXCONFIG=C:\TRANSACTION_SAMPLE\tuxconfig
    TLOGDEVICE=C:\APP_DIR\TLOG
    TLOGNAME=TLOG
    TUXDIR="C:\WLEdir"
    MAXWSCLIENTS=10

*GROUPS
    SYS_GRP
        LMID      = SITE1
```

6 Using Transactions

```
        GRPNO      = 1
        ORA_GRP
        LMID       = SITE1
        GRPNO      = 2

        OPENINFO   = "ORACLE_XA:Oracle_XA+SqlNet=ORCL+Acc=P
        /scott/tiger+SesTm=100+LogDir=.+MaxCur=5"
        OPENINFO   = "ORACLE_XA:Oracle_XA+Acc=P/scott/tiger
        +SesTm=100+LogDir=.+MaxCur=5"
        CLOSEINFO  = ""
        TMSNAME    = "TMS_ORA"

* SERVERS
        DEFAULT:
        RESTART  = Y
        MAXGEN   = 5

        TMSYSEVT
        SRVGRP   = SYS_GRP
        SRVID    = 1

        TMFFNAME
        SRVGRP   = SYS_GRP
        SRVID    = 2
        CLOPT    = "-A -- -N -M"

        TMFFNAME
        SRVGRP   = SYS_GRP
        SRVID    = 3
        CLOPT    = "-A -- -N"

        TMFFNAME
        SRVGRP   = SYS_GRP
        SRVID    = 4
        CLOPT    = "-A -- -F"

        TMIFRSVR
        SRVGRP   = SYS_GRP
        SRVID    = 5

        UNIVT_SERVER
        SRVGRP   = ORA_GRP
        SRVID    = 1
        RESTART  = N

        ISL
        SRVGRP   = SYS_GRP
        SRVID    = 6
        CLOPT    = -A -- -n //MACHINENAME:2500
```

*SERVICES

For information about the transaction log and defining parameters in the Configuration file, see [*Creating a Configuration File*](#) in the WebLogic Enterprise online documentation.

Part III Developing WebLogic Enterprise EJB Applications

Chapter 7. Developing WebLogic Enterprise EJB Applications

Chapter 8. Designing Enterprise JavaBeans for the WebLogic Enterprise System

7 Developing WebLogic Enterprise EJB Applications

This chapter provides a step-by-step tutorial that explains how to create an EJB application that you can build and run in the WebLogic Enterprise environment. The steps described in this chapter use the `statefulSession` EJB sample application provided with the WebLogic Enterprise software. The `statefulSession` example exists in the following location on your system:

Windows NT

```
$TUXDIR\samples\j2ee\ejb\basic\statefulSession
```

UNIX

```
$TUXDIR/samples/j2ee/ejb/basic/statefulSession
```

This topic includes the following sections:

- Overview of the Development Process for WebLogic Enterprise EJB Applications
- The `statefulSession` EJB Sample Application
- Developing EJB Applications
- Building and Deploying EJB Applications
- WebLogic Enterprise EJB Sample Applications

Note the following about using this chapter:

- The steps in this chapter lead you through the process of creating an EJB that can be deployed in the WebLogic Enterprise environment, also known as a deployable EJB. To create a standard EJB -- that is, a portable bean that lacks the extensions specific to the WebLogic Enterprise environment -- see Chapter 8, "Designing Enterprise JavaBeans for the WebLogic Enterprise System."
- For a more thorough understanding of Enterprise JavaBeans, especially with regards to background information and programming considerations, refer to the following resources:
 - Valesky, T. 1999. *Enterprise JavaBeans: Developing Component-Based Distributed Applications*. Addison-Wesley.
 - Sun Microsystems, Inc. Enterprise JavaBeans Specification 1.1.

Overview of the Development Process for WebLogic Enterprise EJB Applications


Table 7-1 outlines the development process for WebLogic Enterprise EJB applications.

Table 7-1 EJB Application Development Process in WebLogic Enterprise

Step	Description
1	Create the EJB.
2	Create the module initializer class.
3	Create the deployment descriptor.
4	Create a standard EJB JAR file.
5	Create the WebLogic extensions to the deployment descriptor DTD.
6	Modify the Deployment Descriptor.
7	Package the components into a deployable EJB JAR file.
8	Configure the EJB application.
9	Create the client application.
10	Start and run the WebLogic Enterprise EJB application.
11	Dynamically manage the EJB deployment.

Figure 7-1 illustrates the process for developing WebLogic Enterprise EJB applications. In this figure, the shaded objects represent entities you need to create.

Figure 7-1 Process for Developing WebLogic Enterprise EJB Applications



The statefulSession EJB Sample Application

The `statefulSession` sample application shows how repeated calls to the same session bean have a persistent state -- the change in the cash account -- that is maintained across all the calls. Notice that neither the client nor the EJB do anything to maintain that state: the container handles it transparently. All the logic for the cash account is encapsulated in the bean, unlike the stateless session sample where all persistence is provided by the client.

The EJB in this sample provides basic trading methods such as buying and selling stocks. Since there are no persistent stores involved in this sample, all the stock data are set in the deployment descriptor of the EJB as environment properties. The container supplies the data to the EJB through the JNDI lookup operation.

This sample provides two types of clients: one is a simple, single-threaded RMI client application, and the other is a multithreaded RMI client application. The `statefulSession` bean sample application implements the classes listed and described in Table 7-2.

Table 7-2 Classes Implemented in the Stateful Session Bean Example


Class	Description
Client	This class: <ul style="list-style-type: none">■ Creates an <code>InitialContext</code> class.■ Creates a trader, and performs repeated buying and selling of shares.■ Shows persistence of state between calls to the <code>TraderBean</code>; the client does not do anything to maintain state between calls.■ Searches the JNDI tree for an appropriate container.

Table 7-2 Classes Implemented in the Stateful Session Bean Example

Class	Description
MultiClient	<p>This class:</p> <ul style="list-style-type: none">■ Creates an InitialContext class.■ Creates a trader, and performs repeated buying and selling of shares.■ Shows calling a stateful session bean using multiple colocated clients: each thread is a trader, and performs repeated buying and selling of shares.■ Shows persistence of state between calls to the TraderBean bean. <p>Like the single-threaded Client bean, the MultiClient bean does not do anything to maintain state between calls.</p>
TraderBean	This bean does not manage any persistence of state between invocations on it. Creating the business methods on this bean is described in the section “Create the Bean’s Implementation Class” on page 7-11.
TradeResult	This bean contains the results of a buy/sell transaction.

Figure 7-2 shows how the stateful session bean example works.

Figure 7-2 Stateful Session Bean Sample Application



Developing EJB Applications

This section describes the following steps for developing an EJB application in the WebLogic Enterprise system:

- Step 1: Create the EJB
- Step 2: Create the Module Initializer Object
- Step 3: Create the Deployment Descriptor

- Step 4: Create a Standard EJB JAR File

Step 1: Create the EJB

The EJB Specification 1.1, published by Sun Microsystems, Inc., describes the different requirements of the EJB writer and the EJB framework; EJBs created for the WebLogic Enterprise environment must conform to those requirements. When you write EJBs, pay close attention to these requirements.

When writing an EJB, you must implement the following:

- The bean's home interface
- The bean's remote interface
- The bean's implementation class, which includes:
 - The business methods for the bean
 - The `ejbCreate`, `ejbPostCreate`, and `ejbRemove` callback methods
 - For session beans, the callbacks defined by the `SessionBean` interface and, optionally, the callbacks on the `SessionSynchronization` interface
 - For entity beans, the callbacks defined on the `EntityBean` interface, and the primary key classes
 - For bean-managed persistence, the `ejbLoad` and `ejbStore` callbacks

Notes: The direct use of threads by Bean Providers is discouraged by the EJB Specification 1.1. This constraint also applies to WebLogic Enterprise server applications -- bean and RMI implementers should not attempt to manage, change properties, start, stop, suspend, or resume a thread or a thread group.

The `ejbc` command, which is provided with the WebLogic Enterprise development software, includes a compliance checker utility that examines the packaged EJBs and determines if the EJBs conform to these requirements.

The subsections that follow provide details on implementing an EJB, using the stateful session EJB sample application as an example.

Create the Bean's Home Interface

Each EJB has a home interface that creates instances of the bean. EJB client applications use the home interface as a means of obtaining a reference, or a handle, to the EJB. The home interface is analogous to a factory object in CORBA. The home interface defines the methods used by client applications to create, remove, and find objects of the corresponding EJB type.

The home interface for the `statefulSession` EJB contains the `create` method, which corresponds to the `ejbCreate` method on the EJB itself. The following code example shows the home interface for the `TraderBean` EJB:

```
package samples.j2ee.ejb.basic.statefulSession;

import java.rmi.RemoteException;
import javax.ejb.*;

/**
 * This interface is the home interface for the TraderBean.java
 */
public interface TraderHome extends EJBHome {
    Trader create(String traderName) throws CreateException, RemoteException;
}
```

Create the Bean's Remote Interface

Each EJB has a well-defined remote interface that defines the EJB callbacks and the business methods that can be invoked by a client. As stated in the EJB Specification 1.1, a client application never directly accesses instances of a bean's class. A client always uses the bean's remote interface to access that bean's instance. The class that implements the bean's remote interface is provided by the EJB container.

The EJB's remote interface does the following:

- Defines the business logic methods of the EJB.
- Supports the methods of the `javax.ejb.EJBObject` interface. These methods allow the client to:
 - Get the EJB's home interface
 - Get the EJB's handle
 - Test if the EJB is identical with another EJB

- Remove the EJB

The following business methods are also defined on the remote interface of the `TraderBean` EJB:

- `buy()`
- `sell()`
- `getBalance()`
- `getTraderName()`

Listing 7-1 shows the remote interface for the `TraderBean` EJB.

Listing 7-1 TraderBean Remote Interface

```
package samples.j2ee.ejb.basic.statefulSession;

import java.rmi.RemoteException;
import javax.ejb.*;

/**
 * The methods in this interface are the public face of TraderBean.
 * The signatures of the methods are identical to those of the EJBean, except
 * that these methods throw a java.rmi.RemoteException.
 */

public interface Trader extends EJBObject {

    public TradeResult buy(String customerName, String stockSymbol, int shares)
        throws ProcessingErrorException, RemoteException;

    public TradeResult sell(String customerName, String stockSymbol, int shares)
        throws ProcessingErrorException, RemoteException;

    public double getBalance()
        throws RemoteException;

    public String getTraderName()
        throws RemoteException;
}
```

Create the Bean's Implementation Class

The bean's implementation class includes the following:

- A declaration of the type of bean being implemented (session or entity).
- The business methods on the bean.
- A specific set of callback methods. The set you implement depends on the type of bean you are creating.

Declaring the Bean Type

After you declare your bean's package name and import classes, you declare what interface your bean implements: session or entity. The following line declares that the `TraderBean` class implements the `SessionBean` interface:

```
public class TraderBean implements SessionBean
```

Implementing the Business Methods on the Bean

The `TraderBean` EJB implements the following business methods:

- `buy()`
Buys shares of a stock for a named customer.
- `sell()`
Sells shares of a stock for a named customer.
- `getBalance()`
Returns the current balance of a trading session.
- `getTraderName()`
Returns the name of the current Trader class.

Implementing the Callback Methods on the Bean

You need to implement the following methods on the bean, which are standard for all beans. These are commonly referred to as callback methods.

- `ejbCreate()`

Corresponds to the `create` method in the home interface `TraderHome.java`. The parameter sets of the `ejbCreate` and `create` methods are identical. When the client calls the `create` method on the bean's home interface, the EJB container allocates an instance of the `TraderBean` EJB and then calls the `ejbCreate` method on the EJB.

- `ejbRemove()`

This method is automatically invoked by the EJB container before it ends the life of the session bean. Invoking this method occurs as a result of a client application invoking the `remove` method on the bean's remote object, or when the EJB container decides to terminate the session object after a timeout. This method is required by the EJB Specification.

- `ejbPassivate()`

This method is invoked by the EJB container before a bean is moved into a passive state, causing any resources used by the bean to be released. This method is required by the EJB Specification, but is not used by the `statefulSession` bean example.

- `ejbActivate()`

This method is invoked by the EJB container when a bean is activated from a passive state, causing any resources required by the bean to be restored. This method is required by the EJB Specification, but is not used by the `statefulSession` bean example.

- `setSessionContext()`

This method sets the associated session context. The WebLogic Enterprise container invokes this method after the EJB has been instantiated. The EJB instance should store the reference to the context object in an instance variable.

ejbCreate Callback Example

The following code example shows the `ejbCreate` method on the `TraderBean` EJB:

```
public void ejbCreate(String traderName) throws CreateException {  
    printTrace("ejbCreate (" + traderName + ")");  
    this.traderName = traderName;  
    this.tradingBalance = 0.0;
```

setSessionContext Callback Example

The following code example shows the `setSessionContext` method on the `TraderBean` EJB, storing the context in the variable `ctx`:

```
public void setSessionContext(SessionContext ctx) {  
    printTrace("setSessionContext called");  
    this.ctx = ctx;  
}
```

Step 2: Create the Module Initializer Object

The module initializer object is optional for EJB applications that run in the WebLogic Enterprise environment. You use it for specifying special requirements for your EJB application, such as custom operations; for example:

- Performing basic module initialization (or EJB JAR file deployment) operations, which may include allocating resources needed by the EJB JAR file.
- Performing basic server application initialization operations, which may include registering homes or factories managed by the server application and allocating resources needed by the server application.
- Performing server process shutdown and cleanup procedures when the server application has finished servicing requests.

Notes: For EJBs, the scope of the module initializer object is at the EJB JAR file level and not of the entire server application, as with the `Server` object and WebLogic Enterprise CORBA applications.

The `statefulSession` EJB sample application does not include a module initializer object.

If you have enabled hot redeployment for the modules in your EJB application, the module initializer object is automatically invoked at appropriate times when the module is deployed or undeployed.

You implement this module initializer object by creating a module initializer class that derives from `com.beasys.Tobj.Server` and by implementing the following two methods on that class:

- `initialize`

The `initialize` method is invoked when the EJB JAR file is loaded (generally when the WebLogic Enterprise server process is booted).

- `release`

The `release` method is invoked when the WebLogic Enterprise server process is shut down or when the EJB JAR file is redeployed to another server process.

In the module initializer object application code, you can also write a public default constructor. You create the module initializer object class from scratch using a text editor.

If you have created a module initializer object, the EJB container parses the WebLogic EJB extensions to the deployment descriptor DTD in each deployed EJB JAR file (specified in the `UBBCONFIG` file) during startup.

The `module-initializer-class-name` element in the WebLogic EJB extensions to the deployment descriptor DTD identifies the module initializer object to be used at server initialization and shutdown or, if you are using hot redeployment, when a module is deployed or undeployed. When the server process is booted or a module is deployed, the EJB container instantiates this module initializer object and invokes its `initialize` method, passing in any startup arguments specified in the `UBBCONFIG` file. When the server process is shut down or a module is undeployed, the EJB container invokes the module initializer object's `release` method.

For information about the `com.beasys.Tobj.Server` base class, see the [API Javadoc](#) in the WebLogic Enterprise online documentation. For more information about hot redeployment, see “*Step 11: Dynamically Manage the EJB Deployment (Hot Redeployment)*” on page 7-37.

Step 3: Create the Deployment Descriptor

The deployment descriptor is an XML file that specifies structural information (for example, the name of the enterprise bean class) about the EJB and declares all the EJB's external dependencies (for example, the names and types of resources that the enterprise bean uses). For complete details on all the elements you can specify in the deployment descriptor, see the [EJB XML Reference](#) in the WebLogic Enterprise online documentation.

The deployment descriptor also ties together the different classes and interfaces, and is used by the `e jbc` command to build the code-generated class files. You can also use the deployment descriptor to specify critical aspects of the EJB's deployment at run time.

You create the deployment descriptor using one of the following methods:

- Using the WebLogic EJB Deployer
- Using the `DDGenerator` command
- Manually, using a text editor

Note that the `DDGenerator` command is an unsupported tool. For information about using the `DDGenerator` command, see the *Release Notes*.

The deployment descriptor you create must:

- Be valid with respect to the Document Type Definition (DTD) documented in the EJB Specification 1.1
- Conform to the semantics rules specified in the DTD comments and elsewhere in the EJB Specification 1.1
- If you are creating the deployment descriptor from scratch, include the following reference to the deployment descriptor DTD at the beginning of the file:

```
<!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise JavaBeans 1.1//EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_2.dtd">
```

The sections that follow describe the elements you must specify in a deployment descriptor for a given type of EJB, using the `statefulSession` as an example, and show the EJB Deployer.

Required Elements for Session Beans

The elements that you need to specify in the deployment descriptor for the stateless session EJB is listed and described in Table 7-3.

Table 7-3 Required Elements for Session Beans

Element	Description	Purpose
ejb-name	EJB's name	Specifies the logical name you assign to each EJB in the EJB JAR file. There is no architected relationship between this name and the JNDI name that the Deployer assigns to the EJB.
ejb-class	EJB's class	Specifies the fully qualified name of the Java class that implements the EJB's business methods.
home	EJB's home interface	Specifies the fully qualified name of the EJB's home interface.
remote	EJB's remote interfaces	Specifies the fully qualified name of the EJB's remote interface.
session entity	EJB's type	The EJB types are <code>session</code> and <code>entity</code> . Use the appropriate <code>session</code> or <code>entity</code> element to declare the EJB's structural information.
session-type	Session bean's state management type	Declares whether the session bean is stateful or stateless.
transaction-type	Session bean's transaction demarcation type	If the EJB is a session bean, declares whether transaction demarcation is performed by the enterprise bean or by the container.

Listing 7-2 shows the deployment descriptor for the `statefulSession` EJB. Line numbers are added to help with the discussion of this deployment descriptor, which follows.

Listing 7-2 Stateful Session Bean Deployment Descriptor

```
1  <!DOCTYPE ejb-jar PUBLIC "-//Sun Microsystems, Inc.//DTD Enterprise
2  JavaBeans 1.1//EN" "http://java.sun.com/j2ee/dtds/ejb-jar_1_2.dtd">
3
4  <ejb-jar>
5    <enterprise-beans>
6      <session>
7        <ejb-name>
```

```
8      statefulSession
9  </ejb-name>
10 <home>
11   samples.j2ee.ejb.basic.statefulSession.TraderHome
12 </home>
13 <remote>
14   samples.j2ee.ejb.basic.statefulSession.Trader
15 </remote>
16 <ejb-class>
17   samples.j2ee.ejb.basic.statefulSession.TraderBean
18 </ejb-class>
19 <!-- Session bean state management type declaration -->
20 <session-type>
21   Stateful
22 </session-type>
23 <!-- Transaction demarcation type declararion -->
24 <transaction-type>
25   Container
26 </transaction-type>
27
28 <!-- Environment entries: Stock symbols, and prices -->
29 <env-entry>
30   <env-entry-name>
31     BEAS
32   </env-entry-name>
33   <env-entry-type>
34     java.lang.Double
35   </env-entry-type>
36   <env-entry-value>
37     10.0
38   </env-entry-value>
39 </env-entry>
40   .
41   .
42   .
43 </session>
44 </enterprise-beans>
45
46 <!-- Assembly description -->
47 <assembly-descriptor>
48
49 <!-- Container transaction attributes -->
50 <container-transaction>
51   <method>
52     <ejb-name>
53       statefulSession
54     </ejb-name>
55
56   <!-- Apply to all the methods (*) of the ejb -->
```

```
57      <method-name>
58          *
59      </method-name>
60  </method>
61
62      <!-- Transaction attributes for the methods -->
63      <trans-attribute>
64          NotSupported
65      </trans-attribute>
66  </container-transaction>
67  </assembly-descriptor>
68 </ejb-jar>
```


In the preceding deployment descriptor, note the following:

- Lines 7-9 identify `statefulSession` as the EJB name.
- Lines 10-15 identify the home and remote interfaces.
- Lines 20-22 identify that the `statefulSession` EJB is stateful.
- Lines 28-42 identify a number of environment entries specific to the `statefulSession` EJB.
- Lines 46-67 identify assembly descriptor information for transaction attributes. For information about EJBs and transactions, see [Using Transactions](#) in the WebLogic Enterprise online documentation.

Using the WebLogic EJB Deployer to Create the Deployment Descriptor

You can use the WebLogic EJB Deployer to create the deployment descriptor for the EJB. Figure 7-3 shows the EJB Deployer main window.

Figure 7-3 WebLogic EJB Deployer Main Window



For information about how to start and use the EJB Deployer, see [Using the WebLogic Enterprise EJB Deployer](#) in the WebLogic online documentation.

Step 4: Create a Standard EJB JAR File

In this step, you create a standard EJB JAR file. A standard EJB JAR file contains an EJB that has been built, but lacks the specific deployment information on any specific system. You typically build a standard EJB with the goal of being able to distribute that EJB to a variety of deployment environments. Typically, the Bean Provider performs

steps 1 through 4, as described in the chapter, and a standard EJB JAR file provides a convenient package that can be handed off to the Application Assembler or Deployer, who may perform steps 5 through 9.

The input to the standard EJB JAR file is typically:

- Compiled class files for the EJB's implementation class, the home interface, and remote interface
- The deployment descriptor

You can create a standard EJB JAR file using the `jar` command.

Building and Deploying EJB Applications

This section describes the steps to develop an EJB application for the WebLogic Enterprise system:

- Step 5: Create the WebLogic EJB Extensions to the Deployment Descriptor DTD
- Step 6: Modify the Deployment Descriptor
- Step 7: Package the Components Into a Deployable EJB JAR File
- Step 8: Configure the EJB Application
- Step 9: Create the Client Application
- Step 10: Start and Run the WebLogic Enterprise EJB Application

Step 5: Create the WebLogic EJB Extensions to the Deployment Descriptor DTD

For an EJB application to be deployable in the WebLogic Enterprise environment, you need to create a file containing the WebLogic EJB extensions to the deployment descriptor DTD. This file specifies the following run time and configuration information for the EJB application:

- Custom application startup and shutdown properties
- Registration of the application's home interfaces
- Persistence information

For complete details on all the elements you can specify in the WebLogic EJB extensions to the deployment descriptor DTD, see the [EJB XML Reference](#) in the WebLogic Enterprise online documentation.

Creating the WebLogic EJB Extensions to the Deployment Descriptor DTD

You can create the file containing the WebLogic EJB extensions to the deployment descriptor DTD using one of the following methods:

- Using the WebLogic EJB Deployer
- By hand, in a common text editor

Specifying the WebLogic EJB Extensions DTD

The file that includes the WebLogic EJB extensions to the deployment descriptor DTD must specify the following DTD reference at the beginning of the file:

```
<!DOCTYPE weblogic-ejb-extensions SYSTEM "weblogic-ejb-extensions.dtd" >
```

Registering Names for the EJB Home Classes

A name for the EJB home class must be registered in the global WebLogic Enterprise JNDI namespace. This allows Java clients to perform a lookup on the JNDI name for the EJB home, even across WebLogic Enterprise domains, and gain access to the object. The name for the EJB home class can be different than the `<ejb-name>`

element specified in the standard EJB XML. The `<ejb-name>` in the standard deployment descriptor must be unique only among the names of the EJBs in the same EJB JAR file. However, the JNDI name must be unique among all global home or factory names in a WebLogic Enterprise domain; this includes EJB homes, CORBA factories, and RMI-named objects.

Example

Listing 7-3 is from the file `weblogic-ejb-extensions.xml`, which specifies the WebLogic extensions to the deployment descriptor DTD for the stateful session bean example. Line numbers are used to aid in the brief discussion that follows.

Listing 7-3 Specifying the Name of the EJB Home Class

```
1  <weblogic-ejb-extensions>
2    <weblogic-version>
3      WebLogic Enterprise Server 5.0
4    </weblogic-version>
5    <weblogic-enterprise-bean>
6      <ejb-name>
7        statefulSession
8      </ejb-name>
9      <weblogic-deployment-params>
10        <jndi-name>
11          statefulSession.TraderHome
12        </jndi-name>
13        .
14        .
15        .
16      </weblogic-deployment-params>
17  </weblogic-ejb-extensions>
```

In the preceding WebLogic EJB extensions to the deployment descriptor DTD, note the following lines:

- Lines 9 through 16 contain specific deployment parameters for metrics like the size of the bean pool and the size of the cache.
- Lines 10 through 12 specify the name of the EJB that is registered with JNDI; this is the name on which the client application performs a `lookup` invocation.

Specifying Persistence Information

The WebLogic Enterprise EJB container provides container-managed persistence. The code for implementing the persistence is generated by the `e jbc` command based on the deployment descriptors. The persistence store can be a flat file or it can be a database managed with a JDBC connection pool. For the EJB state to fully cooperate in a WebLogic Enterprise global transaction, configure the EJB to use the JDBC-managed database store provided in WebLogic Enterprise. Use file-based persistence only during development and prototyping.

The standard deployment descriptor created by the Bean Provider normally specifies:

- The fields in the EJB that are to be persistent, via the `cmp-field` element
- For entity beans, information about the primary key

However, you, as the deployer, need to specify additional information for mapping an EJB to its persistent store via the WebLogic EJB extensions to the deployment descriptor DTD.

File-based Persistence

Listing 7-4 shows the WebLogic EJB extensions to the deployment descriptor DTD for specifying file-based persistence.

Listing 7-4 File-based Persistence Elements

```
<!--
Persistence store descriptor. Specifies what type of persistence store
EJB container should use to store state of bean.
-->
<!ELEMENT persistence-store-descriptor (description?,
(persistence-store-file |
persistence-store-jdbc)?)>
<!--
Persistence store using file. Bean is serialized to a file.
Mainly used to store state of Stateful Session Beans.
-->
<!ELEMENT persistence-store-file (description?,
persistence-store-directory-root
?)>
<!--
Root directory on File system for storing files per bean.
```

```
-->
<!ELEMENT persistence-store-directory-root (#PCDATA)>
```

The information supplied for the `persistence-store-directory-root` element is used by the EJB container to store all instances of the EJB, with the `ejb-name` element converted to a directory name.

Database-stored Persistence

Listing 7-5 shows the WebLogic EJB extensions to the deployment descriptor DTD for specifying a JDBC connection for database-stored persistence.

Listing 7-5 Database-stored Persistence Elements

```
<!--
Persistence store is any RDBMS. JDBC driver is used to talk to database.
Required for CMP.
-->
<!ELEMENT persistence-store-jdbc (description?, pool-name, user?, password?,
driver-url?, driver-class-name?, table-name, attribute-map,
finder-descriptor*)>

<!-- Required for CMP -->
<!ELEMENT pool-name (#PCDATA)>

<!-- Ignored in WebLogic Enterprise Server as this is part of connection
pool
setup at startup -->
<!ELEMENT user (#PCDATA)>

<!-- Ignored in WebLogic Enterprise Server as this is part of connection
pool
setup at startup -->
<!ELEMENT password (#PCDATA)>

<!-- Ignored in WebLogic Enterprise Server as this is part of connection
pool
setup at startup -->
<!ELEMENT driver-url (#PCDATA)>

<!-- Ignored in WebLogic Enterprise Server as this is part of connection
pool
setup at startup -->
<!ELEMENT driver-class-name (#PCDATA)>
```

```
<!-- Required for CMP -->
<!ELEMENT table-name (#PCDATA)>

<!-- Required for CMP -->
<!ELEMENT attribute-map (description?, attribute-map-entry+)>

<!-- Required for CMP -->
<!ELEMENT attribute-map-entry (bean-field-name, table-column-name)>

<!-- Required for CMP -->
<!ELEMENT bean-field-name (#PCDATA)>

<!-- Required for CMP -->
<!ELEMENT table-column-name (#PCDATA)>

<!-- Required for CMP -->
<!ELEMENT finder-descriptor (description?, method?, query-grammar?)>

<!-- Required for CMP -->
<!ELEMENT query-grammar (#PCDATA)>
```

The EJB instances are stored in a database that has been previously set up with a JDBC connection pool, which is identified by the `pool-name` element. The `table-name` and `attribute-map` elements map the EJB fields to the appropriate table columns in the database.

Finder descriptors are the WebLogic Enterprise implementation of the EJB `find` methods. The `finder-descriptor` elements are pairs of method signatures and expressions. You specify a method signature in the `EJBHome` interface, and you specify the method's expression in the deployment descriptor via the `query-grammar` element. The finder methods return an enumeration of EJBs. For more information about specifying finder descriptors, see the [EJB XML Reference](#) in the WebLogic Enterprise online documentation.

Example

Listing 7-6 is from the file `weblogic-ejb-extensions.xml`, which specifies the WebLogic extensions to the deployment descriptor DTD for the stateful session bean example to show specifying the persistence information.

Listing 7-6 Persistence Directory Root

```
<persistence-store-descriptor>
  <persistence-store-file>
    <persistence-store-directory-root>
      c:\mystore
    </persistence-store-directory-root>
  </persistence-store-file>
</persistence-store-descriptor>
```

Specifying the Module Initializer Class

If your EJB application uses a module initializer class, as explained in the section “Step 2: Create the Module Initializer Object” on page 7-13, you need to specify that class among the XML elements for startup and shutdown procedures in the WebLogic EJB extensions to the deployment descriptor DTD. The WebLogic Enterprise EJB container parses the XML at run time and performs the startup and shutdown processing.

Note: The `statefulSession` EJB sample application does not include a module initializer object.

Step 6: Modify the Deployment Descriptor

The Bean Provider specifies some initial deployment information in the deployment descriptor. The deployer typically needs to add to or modify that information, such as shown in Table 7-4.

Table 7-4 Deployment Descriptor Fields Modified By the Deployer

Field	Description
The EJB’s name	You may change the enterprise bean’s name defined in the <code>ejb-name</code> element.
Values of environment entries	You may change existing values or define new values for the environment properties.
Description fields	You may change existing or create new description elements.

Table 7-4 Deployment Descriptor Fields Modified By the Deployer (Continued)

Field	Description
Binding of enterprise bean references	You may link an enterprise bean reference to another enterprise bean in the EJB JAR file. You create the link by adding the <code>ejb-link</code> element to the referencing bean.
Security roles	You may define one or more security roles. The security roles define the recommended security roles for the clients of the enterprise beans. You define the security roles using the <code>security-role</code> elements. For more information about EJB security, see Using Security in the WebLogic Enterprise online documentation.
Method permissions	You may define method permissions, which are binary relationships between the security roles and the methods of the remote and home interfaces of the EJBs. You define method permissions using the <code>method-permission</code> elements.
Linking of security role references	If you define security roles in the deployment descriptor, you must link the security role references declared by the Bean Provider to the security roles. You define these links using the <code>role-link</code> element. For more information about EJB security, see Using Security in the WebLogic Enterprise online documentation.

Table 7-4 Deployment Descriptor Fields Modified By the Deployer (Continued)

Field	Description
Changing persistent storage information, if necessary	<p>The deployer can change the type of persistent storage used by a bean. If the <code>persistentStoreType</code> is <code>file</code>, the serialized files are created in this directory. The default file is <code>/pstore/bean_name.dat</code>, where the directory <code>pstore</code> represents the directory from which the WebLogic Enterprise application was started, and <code>bean_name</code> is the fully qualified name of the EJB with underscores (<code>_</code>) replacing the periods (<code>.</code>) in the name.</p> <p>If the <code>persistentStoreType</code> is <code>jdbc</code>, the container looks for additional values to determine the appropriate values for the JDBC connection. Note that if the bean's persistence is stored in a database via a JDBC connection, the system administrator needs to add this information to the <code>UBBCONFIG</code> file as well. For more information, see Using the JDBC Drivers in the WebLogic Enterprise online documentation.</p> <p>Note that persistence information is specified in the WebLogic EJB extensions to the deployment descriptor DTD file, as described in the section "Specifying Persistence Information" on page 7-23.</p>

However, if the Bean Provider, the Application Assembler, and the Deployer are the same person, all the information shown in the preceding table may have been specified already in the deployment descriptor step described in "Step 3: Create the Deployment Descriptor" on page 7-14.

Step 7: Package the Components Into a Deployable EJB JAR File

In this step, you package the deployment descriptor, the compiled files for the EJB classes, and any additional required classes into a deployable EJB JAR file. You can package multiple beans together, provided that there is a deployment descriptor for each bean.

You can use the WebLogic Enterprise `ejbc` command to create the deployable EJB JAR file. The `ejbc` command performs the following steps:

1. Parses the standard EJB deployment descriptor and WebLogic Enterprise extended deployment descriptor XML files.
2. Checks the deployment descriptors for semantic consistency, and writes any inconsistencies to standard output.
3. Generates the wrapper Java classes and compiles them. This is performed for each EJB in the deployment descriptor.
4. Packages the XML deployment descriptors and the generated class files into a deployable EJB JAR file.

If you have multiple bean packages meant to be assembled as a deployable unit, the bean packages must be specified in a single deployment descriptor.

The following command line builds the deployable EJB JAR file for the `statefulSession` bean example:

```
java weblogic.ejbc -validate -i ejb-jar.xml -x weblogic-ejb-extensions
ejb-jar-file
```

In the preceding command line:

- The `-i` option specifies the name of the deployment descriptor, `ejb-jar.xml`.
- The `-x` option specifies the name of the WebLogic EJB extensions to the deployment descriptor DTD, `weblogic-ejb-extensions`.
- The file `ejb-jar-file` is the name of the EJB JAR file.

For more information about the `ejbc` command, see [*Commands, System Processes, and MIB Reference*](#) in the WebLogic Enterprise online documentation. Note that using the `-validate` option is recommended.

Step 8: Configure the EJB Application

Because the WebLogic Enterprise software offers great flexibility and many options to application designers and programmers, no two applications are alike. An application, for example, may be small and simple (a single client and server running on one machine) or complex enough to handle transactions among thousands of client and server applications. For this reason, for every WebLogic Enterprise EJB

application being managed, the system administrator must provide a configuration file that defines and manages the components (for example, domains, server applications, client applications, and modules) of that application.

When system administrators create a configuration file, they are describing the WLE application using a set of parameters that the WLE software interprets to create a runnable version of the application. During the setup phase of administration, the system administrator's job is to create a configuration file. The configuration file contains the sections listed in Table 7-5.

Table 7-5 Configuration File Sections

Sections in the Configuration File	Description
RESOURCES	Defines defaults (for example, user access and the main administration machine) for the WebLogic Enterprise CORBA application.
MACHINES	Defines hardware-specific information about each machine running in the WebLogic Enterprise CORBA application.
GROUPS	Defines logical groupings of server applications or CORBA interfaces.
SERVERS	Defines the server application processes (for example, the Transaction Manager) used in the WebLogic Enterprise CORBA application.
SERVICES	Defines parameters for services provided by the WebLogic Enterprise application.
MODULES	Defines information about the EJB application process.
JDBCONNPOOLS	Describes the pooling of JDBC connections for Java servers.

When creating a configuration file for an EJB application meant to be run in a WebLogic Enterprise domain, note the following:

- The WebLogic Enterprise EJB container runs in a JavaServer, which you configure in the SERVERS section.
- You configure the EJB application server process using the `ejb` keyword in the MODULES section. Important: the server group and server ID (specified by the

SRVGRP and SRVID keywords, respectively) for the EJB server process must match the server group and server ID for the JavaServer that runs the EJB container.

- You use the FILE parameter in the configuration information for your EJB application to specify the EJB JAR file.
- For complete details about the UBBCONFIG file, see *Creating a Configuration File* in the WebLogic Enterprise online documentation.

The following section shows an example configuration file used for an EJB application.

Example Configuration File

Listing 7-7 shows a configuration file for the statefulSession bean example. Note in the SERVERS section how the server group and server ID for the JavaServer, in which the EJB container runs, is the same as for the EJB application configured in the MODULES section.

Listing 7-7 Sample UBBCONFIG File

```
#-----
# D:\BEA\WLE51\test\ejb\basic\statefulSession\ubbconfig.nt
# Generated for basic\statefulSession EJB Sample
#-----
*RESOURCES
  IPCKEY      55432
  DOMAINID   ejbsample
  MASTER      SITE1
  MODEL       SHM
  LDBAL       N
#-----
*MACHINES
  "HOWE"
    LMID        = SITE1
    APPDIR     = "D:\BEA\WLE51\test\ejb\basic\statefulSession"
    TUXCONFIG  = "D:\BEA\WLE51\test\ejb\basic\statefulSession\tuxconfig"
    TUXDIR     = "d:\BEA\WLE51\m3"
    MAXWSCLIENTS = 10
#-----
*GROUPS
  SYS_GRP
    LMID      = SITE1
```

```
        GRPNO      = 1
APP_GRP
        LMID      = SITE1
        GRPNO      = 2
#-----
*SERVERS
    DEFAULT:
        RESTART   = Y
        MAXGEN   = 5
    TMSYSEVT
        SRVGRP   = SYS_GRP
        SRVID    = 1
    TMFFNAME
        SRVGRP   = SYS_GRP
        SRVID    = 2
        CLOPT    = "-A -- -N -M"
    TMFFNAME
        SRVGRP   = SYS_GRP
        SRVID    = 3
        CLOPT    = "-A -- -N"
    TMFFNAME
        SRVGRP   = SYS_GRP
        SRVID    = 4
        CLOPT    = "-A -- -F"
    JavaServer
        SRVGRP   = APP_GRP
        SRVID    = 5
        CLOPT    = "-A "
        RESTART  = N
    ISL
        SRVGRP   = SYS_GRP
        SRVID    = 6
        CLOPT    = "-A -- -n //HOWE:7001"
#-----
*MODULES
    ejb
        SRVGRP = APP_GRP
        SRVID  = 5
        FILE   =
"D:\BEA\WLE51\test\ejb\basic\statefulSession\ejb_basic_statefulSession.jar"
*SERVICES
```

Compiling the Configuration File

There are two forms of the configuration file:

- An ASCII version of the file, created and modified with any editor. Throughout the WebLogic Enterprise documentation, the ASCII version of the configuration file is referred to as the `UBBCONFIG` file. The configuration file may, in fact, be given any file name.
- The `TUXCONFIG` file, a binary version of the `UBBCONFIG` file created using the `tmloadcf` command. When the `tmloadcf` command is executed, the environment variable `TUXCONFIG` must be set to the name and directory location of the `TUXCONFIG` file. The `tmloadcf` command converts the configuration file to binary form and writes it to the location specified in the command.

For more information about the `tmloadcf` command, see [Commands, System Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.

Bean Passivation Behavior -- the EJB Cache

The WebLogic Enterprise EJB container has the ability to cache beans across method invocations as well as across transactions. This significantly reduces the frequency of beans being passivated, thereby providing the performance improvement. EJB caching is enabled by default for stateful beans.

For more information about EJB caching, see the section “Controlling the Passivation of Beans -- the EJB Cache” on page 8-9.

Step 9: Create the Client Application

When you create a client application that can invoke methods on a session EJB deployed in the WebLogic Enterprise environment, you need to include code that does the following:

- Creates an `InitialContext` object
- Uses the `lookup` method on the `InitialContext` object to obtain a reference to the the EJB’s home interface
- Invokes the `create` method on the EJB’s home interface to create an instance of the EJB

- Invokes the business methods on the EJB
- Destroys the instance of the bean

The sections that follow show fragments from the `statefulSession` EJB application to show the basic building blocks of an RMI client application that invokes an EJB in the WebLogic Enterprise environment. For complete details on creating an RMI client application, see [Using RMI in a WebLogic Enterprise Environment](#) in the WebLogic Enterprise online documentation.

Creating an InitialContext Object

Each WebLogic Enterprise EJB client application needs to create an `InitialContext` object to store information about the EJB application and the WebLogic Enterprise domain so that the client application can run. The `InitialContext` object is typically created with the following data, which are passed as parameters to the constructor of the `InitialContext` object:

- The address of the entry points to the WebLogic Enterprise environment
- The name of the factory to access the WebLogic Enterprise domain and global naming service
- Security information, such as the type of authentication, the security principal name, credentials, roles, and passwords
- The URL for remote class loading

The `statefulSession` client application implements a method named `newInitialContext`, in which the `InitialContext` object is created as a hash table. This hash table specifies `env` as

`com.beasys.com.jndi.WLEInitialContextFactory`. After the context is created, the client application has access to bean homes in the WebLogic Enterprise domain using WebLogic Enterprise as the name service provider.

The `newInitialContext` method is shown in the following code fragment:

```
static public Context newInitialContext() throws Exception {  
    Hashtable env = new Hashtable();  
  
    // specify an IIOP Listener/Handler for the desired WLE target domain  
    env.put(Context.PROVIDER_URL, url);  
  
    // Name of the factory to access WLE domain and global naming service.  
    env.put(Context.INITIAL_CONTEXT_FACTORY,
```

```
    "com.beasys.jndi.WLEInitialContextFactory");
/* Security style: strong for SSL, simple for Tuxedo, and none
 * for no authentication at all. If no value is specified then, Tuxedo
 * style authentication is attempted.
 */
env.put(Context.SECURITY_AUTHENTICATION, "simple");
if (user != null) {
    printTrace("user: " + user);

    /* Specifies the identity of the principal for authenticating the caller
     * to the WLE domain
    */
    env.put(Context.SECURITY_PRINCIPAL, user);
    if (password == null) {
        password = "";
    }

    // A string password is used for Tuxedo style authentication.
    env.put(Context.SECURITY_CREDENTIALS, password);
} else {
    // User id is null.
    env.put(Context.SECURITY_AUTHENTICATION, "none");
}
return new InitialContext(env);
}
```

The RMI client application in the `statefulSession` bean example invokes this `newInitialContext` method to create its `InitialContext` object, on which the application can then make the appropriate invocations, as shown in the following statement:

```
Context ctx = newInitialContext();
```

The `newInitialContext` method also creates the `InitialContext` object with specific information related to WebLogic Enterprise security. For information about the security data in the `InitialContext` object, and other client application requirements related to WebLogic Enterprise security, see [Using Security](#) in the WebLogic Enterprise online documentation.

Obtaining a Reference to the EJB's Home Interface

The following code fragment, from the `statefulSession` client application, shows the following:

1. Looking up the name of the `TraderBean`'s home interface, `TraderHome`.
2. Using the `PortableRemoteObject` to narrow the reference to type `Trader`.

```
Object objref = ctx.lookup("statefulSession.TraderHome");
TraderHome brokerage = (TraderHome) PortableRemoteObject.narrow(objref,
                                                               TraderHome.class);
```

Creating an Instance of the EJB

Before a client application can invoke business methods on a bean, the bean needs to be instantiated. EJBs are instantiated when the client application invokes the `create` method on the EJB's home interface.

The following code fragment, from the `statefulSession` bean example, shows invoking the `create` method on the `TraderHome` class, which causes the `TraderBean` and its remote interface to be instantiated:

```
Trader trader = brokerage.create("Terry");
```

Destroying the Instance of the Bean

After a client application is finished with the bean, it is good programming practice to include an invocation to the bean's `remove` method, as in the following example:

```
trader.remove();
```

Step 10: Start and Run the WebLogic Enterprise EJB Application

Use the `tmboot` command to start the server processes in your WebLogic Enterprise EJB application. The EJB application is usually booted from the machine designated as the `MASTER` in the `RESOURCES` section of the `UBBCONFIG` file.

For the `tmboot` command to find executables, the WebLogic Enterprise system processes must be located in `$TUXDIR/bin`. Server applications should be in `APPDIR`, as specified in the configuration file.

When booting server applications, the `tmboot` command uses the `CLOPT`, `SEQUENCE`, `SRVGRP`, `SRVID`, and `MIN` parameters from the configuration file. Server applications are booted in the order in which they appear in the configuration file.

For details about starting and running the `statefulSession` bean example, see *Guide to the EJB Sample Applications*.

For more information about using the `tmboot` command, see [Commands, System Processes, and MIB Reference](#) in the WebLogic Enterprise online documentation.

Step 11: Dynamically Manage the EJB Deployment (Hot Redeployment)

This step is optional. The WebLogic Enterprise system provides a means, sometimes referred to as **hot redeployment**, to dynamically make the following changes to the modules in a running EJB application:

- Add new modules
- Remove existing modules
- Update deployed modules

Note: A module in an EJB application is a unit of deployment that is more than simply the classes in an EJB JAR file. For example, a module can consist of support libraries.

With WebLogic Enterprise hot redeployment:

- You can change modules used in an EJB application without shutting the application down or starting or restarting JavaServers.
- Clients connected to other EJB modules in the WebLogic Enterprise domain that are not the target of a hot redeployment have no interruption in service.

- Client applications of a module affected by hot redeployment will experience an interruption of service, but the hot redeployment does not require the shutting down or starting of additional server processes.
- Hot redeployment is disabled by default.

To use hot redeployment to add, change, or remove a module in a running EJB application:

- Make sure that the module's classes are not present in the default CLASSPATH.
- In the UBBCONFIG file, enable hot deployment for the JavaServer by specifying the `-Dwle.dynamic` option in the CLOPT attribute.
- While the EJB application is running, use the following tmadmin commands to, respectively, add, remove, or change modules:
 - `addmodule`
 - `removemode`le
 - `changemode`le

Notes: The TMIB `T_MODULE` class has been modified to represent modules that can and cannot be modified.

If you attempt to redeploy a module in a JavaServer for which hot redeployment has not been enabled, no change to the module will occur.

For more information about using hot redeployment, see the following:

- For information about specifying the JavaServer in the UBBCONFIG file, see the topic “Starting JavaServer” in [Creating a Configuration File](#).
- For information about specifying the `-Dwle.dynamic` option in the CLOPT attribute, see the section “Using Server Command-line Options” in [Creating a Configuration File](#).
- For more information about the `addmodule`, `removemode`le, and `changemode`le commands, see the `tmadmin` command in Section 1 of the [BEA Tuxedo Reference Manual](#).

WebLogic Enterprise EJB Sample Applications

WebLogic Enterprise provides the sample applications described in the following list. For more information about these samples, how they work, and how you can build and run them, see [Guide to the EJB Sample Applications](#).

- `samples.j2ee.ejb.basic.statelessSession`
Shows a stateless session bean in which the client application must maintain any state across invocations to that bean.
- `samples.j2ee.ejb.basic.statefulSession`
Shows a session bean that uses stateful persistence.
- `samples.j2ee.ejb.sequence.jdbc`
Shows an entity bean that automatically generates its primary key by calling directly to a database using a connection pool and JDBC.
- `samples.j2ee.ejb.sequence.oracle`
Shows an entity bean that automatically generates its primary key by calling directly to a database using a connection pool and an Oracle database.
- `samples.j2ee.ejb.subclass.parent`
Shows a stateless session bean called `ParentBean` that is the parent class for another bean, `ChildBean`.
- `samples.j2ee.ejb.subclass.child`
Shows a stateless session bean called `ChildBean` that inherits methods from a `ParentBean`, and also shows one bean calling another bean.

8 Designing Enterprise JavaBeans for the WebLogic Enterprise System

This topic includes the following sections:

- Designing EJB Applications for the WebLogic Enterprise System
- EJBs and Persistence

The information in this chapter supplements the Sun Microsystems, Inc. Enterprise JavaBeans Specification 1.1.

Designing EJB Applications for the WebLogic Enterprise System

The WebLogic Enterprise software complies with the EJB Specification 1.1. However, to design EJB applications that take advantage of the WebLogic Enterprise architecture, you need to follow certain design rules and patterns. This section describes these design considerations with respect to the following perspectives on the WebLogic Enterprise EJB environment:

- The client application programmer's view
- The bean programmer's view

The Client Application Programmer's View

Client application programmers using EJBs have a uniform development model they can use for beans regardless of whether the beans are local or remote. For each EJB, client programmers have access to the information in Table 8-1.

Table 8-1 EJB Information that Client Programmers Can Access

Information	Description
The bean's home interface	Each EJB has a home interface (factory) that creates instances of the bean. The home interface defines the methods used by client programmers to create, remove, and find objects of the corresponding EJB type. To find a reference to a particular bean home interface, client applications should use the <code>lookup</code> method of the <code>InitialContext</code> object with the <code>PortableRemoteObject</code> class. In the WebLogic Enterprise system, home interfaces are available, and can be advertised, across domains.
The bean's remote interface	Each EJB has a well-defined remote interface that defines the business methods that can be invoked by a client.

Table 8-1 EJB Information that Client Programmers Can Access (Continued)

Information	Description
The object identity	Each EJB instance lives in a home and has a unique identity within its home. The identity of a session bean is generated by the EJB container and is not exposed to the client. The Bean Provider generates the identity of an entity bean (the primary key) and a client can retrieve the primary key from the entity object reference.
The bean's metadata interface	The metadata interface allows clients (typically, application assembly tools) to discover the metadata information about the bean.
The object handle	The handle identifies the object in a portable way. The handle can be serialized. Having a serialized handle lets you store the handle and then use it at a later time, possibly in a different process or in a different system, or by another bean or object. Handles are more useful with entity beans than with session beans.

To access any EJB, a client application needs to obtain a reference from the bean's home interface (factory); however, because the home is also an object, the client needs also to obtain a reference to it. Getting a home reference to the client application presents a bootstrapping problem. However, when you register home references with a directory service -- namely, JNDI -- client applications have a means to obtain a reference to the bean's home, even across WebLogic Enterprise domains. This is exactly what the WebLogic Enterprise EJB container provides.

From the WebLogic Enterprise EJB container perspective, client applications (including JSP and servlets acting as clients) are nontrusted entities that require authentication. They typically require a network connection to access the WebLogic Enterprise EJB container because they run on nontrusted machines.

How to set up this network connection is another bootstrapping problem. This is solved in WebLogic Enterprise by providing a JNDI implementation that runs within the EJB container trusted environment, and by establishing the network connection when the JNDI initial context is created. The parameters required for the initialization are Java properties (name/value pairs) passed as arguments to the constructor of the `InitialContext` object.

A WebLogic Enterprise client application can set the properties shown in Table 8-2.

Table 8-2 InitialContext Object Parameters

Property Name	Values	Description
WLEContext. INITIAL_CONTEXT_FACTORY	com.beasys.jndi. WLEInitialContextFactory	Specifies the WebLogic Enterprise JNDI.
WLEContext.PROVIDER_URL	corbaloc://<host:port> or corbalocs://<host:port>	Defines the address of the entry points to the WebLogic Enterprise environment. The identifier <code>corbaloc</code> indicates that the protocol is WebLogic RMI on IIOP. The identifier <code>corbalocs</code> indicates that the protocol is WebLogic RMI on IIOP with SSL.
WLEContext. SECURITY_AUTHENTICATION	none simple strong	Defines the type of authentication: <ul style="list-style-type: none"> ■ <code>none</code> means no authentication (this is the default). ■ <code>simple</code> means WebLogic Enterprise authentication. ■ <code>strong</code> means SSL authentication (certificate-based).
WLEContext. SECURITY_PRINCIPAL	<Principal Identifier>	Specifies the security principal name. For more information, see Using Security in the WebLogic Enterprise online documentation.
WLEContext. SECURITY_CREDENTIALS	<SSL credentials> or <User Password>	Specifies the credentials when authentication is <code>strong</code> , or the user password when authentication is <code>simple</code> .
WLEContext.CLIENT_NAME	<Security role>	Specifies the security role name used by <code>simple</code> authentication. For more information, see Using Security in the WebLogic Enterprise online documentation.
WLEContext. SYSTEM_PASSWORD	<password>	Specifies the system password if the <code>simple</code> authentication is in effect.
WLEContext.CODEBASE	<url>	Specifies the URL for remote class loading.

The client application is implicitly associated with the security context specified when the `InitialContext` object is created. To specify a new security context -- for example, to invoke objects in a different WebLogic Enterprise domain -- the client application needs to close the current context and establish a new context with the new security attributes. To find a reference to a particular bean home interface, client applications should use the `lookup` method of the `InitialContext` object with the `PortableRemoteObject` class. Client applications can also use the `lookup` method to obtain a reference to the `UserTransaction` object. WebLogic Enterprise client applications cannot modify the WebLogic Enterprise JNDI naming context; that is, client applications can perform only `lookup` operations on this context.

The client can use the bean home interface to find or create session or entity bean instances. The `create` method provided by the EJB home interface creates the requested EJB and returns a reference to it. The client uses the reference for as long as it needs, and when it finishes, can invalidate the reference (and eventually the EJB instance) by invoking the `ejbRemove` method on the EJB instance. Between these method invocations, the client can invoke (optionally within a transaction) any of the business methods provided by the EJB.

The WebLogic Enterprise system complies with the EJB Specification 1.1, and client programmers must be aware of the subtle programming differences provided by the different bean types (refer to the EJB Specification for more details).

Also note that the WebLogic RMI on IIOP protocol is not currently supported for applets running on a Web browser.

The EJB Programmer's View

As an EJB programmer, or **Bean Provider**, as identified in the EJB Specification, you must follow the conventions and programming restrictions established in the EJB Specification 1.1 for the different EJB types. The following are the principal design considerations to take into account when implementing beans with the WebLogic Enterprise EJB environment:

- Choosing Between Session and Entity Beans
- Server Startup
- Home Interface Registration
- Bean Activation and Passivation

- EJBs as Client Applications

The sections that follow discuss each of these considerations in detail.

Choosing Between Session and Entity Beans

When to use one bean type or another depends upon the design pattern that bean programmers want to use. There are a few commonly used rules:

- Stateless session beans provide a capability similar to the service model provided by the Tuxedo system. They are highly recommended for short interactions with the business data when there is no need to retain state. Therefore, they do not need special operations to activate or deactivate their state. EJB containers can freely pool instances, allocate instances as needed, and apply load balancing strategies to distribute the load across different servers.
- Stateful session beans are recommended when it is necessary to retain in-memory state across multiple method invocations made by the client. These beans are more expensive than stateless session beans because they allocate and exclusively reserve resources during the private conversation with the client application.

When developing a stateful session bean, you must implement the `e jbPassivate` and `e jbActivate` methods in such a way that resources like JDBC connections and network connections are handled properly. While the EJB container is responsible for saving the conversational state in a portable way and for reconstructing that state during activation, some precautions must be taken by the Bean Provider to ensure that the EJB container handles state correctly. (See section 6.4.1 in the EJB Specification 1.1 for more information.) You can also decide if the bean's state needs to be synchronized when the bean is involved in a distributed transaction.

- Entity beans are recommended when it is necessary to associate a bean instance with a particular application-defined identity, which is similar to the CORBA model, and the bean's state must be persistent (that is, the state cannot be lost). Entity beans cannot use the `SessionSynchronization` interface to synchronize with the starting and stopping of a transaction.

Entity beans can be used in many ways; for example, to implement a persistent variant of CORBA objects or to provide an object representation of entities stored in a database. You must be careful when you use entity beans to model objects stored in a database, because these beans could introduce inefficiencies, such as having most of the business logic on the client application rather than in

the server application. Moving the business logic to the server application reduces the number of invocations needed to perform a particular business transaction.

Server Startup

The WebLogic Enterprise EJB container gives you the flexibility to specify an object that is invoked by the EJB container when a WebLogic Enterprise server process loads an EJB JAR file. This object is an instance of a class that implements the `Server` interface, and is thus referred to as a **module initializer object**. Implementing the module initializer object is described in the section “Step 2: Create the Module Initializer Object” on page 7-13.

You can use the module initializer object to perform specialized initialization for some objects, such as instantiating RMI objects. In WebLogic Enterprise EJB applications, you can specify this class in the `<module-initializer-class>` element in the WebLogic EJB extensions to the deployment descriptor DTD, which is a special deployment descriptor that you create along with the standard deployment descriptor.

Home Interface Registration

When an EJB JAR file is deployed, the WebLogic Enterprise container recognizes the EJB home interfaces (factories) and automatically registers them within the WebLogic Enterprise JNDI context. The information about the home interfaces is retrieved from the deployment descriptor.

Bean Activation and Passivation

EJB containers have complete control of the passivation of EJBs. This allows the container to pool instances of a bean and to decide when an instance can be passivated (or removed from the pool) to provide better use of system resources.

As a bean programmer, you cannot make any assumptions about when a bean object is passivated. This passivation can happen at any time. This is particularly important when the bean accesses a database via cursors, because these cursors could become invalid after the passivation; the EJB container can reactivate the bean in another server process.

The WebLogic Enterprise EJB container currently follows a passivation model that is similar to the model used by the WebLogic Enterprise TP Framework for CORBA applications. If resources are scarce, the WebLogic Enterprise EJB container may passivate an object at any time. When the bean is reactivated, it may be reactivated in the same server process or in another server process in the same group.

If a client application creates or invokes a stateful or entity bean within a transaction, the bean will never be passivated while it is participating in the transaction. If the client invocation is nontransactional, the bean may be passivated at the end of the method invocation.

For concurrency control, the WebLogic Enterprise system applies the following rules:

- For entity beans, for a given primary key, there is only one active instance of the bean at one time. This constraint is compatible with the activation policy provided by the WebLogic Enterprise TP Framework.
- Although the client application can issue concurrent invocations on a bean, the WebLogic Enterprise EJB container queues concurrent invocations on a bean so those invocations are performed one at a time. The WebLogic Enterprise system enforces this rule by running each active object on its own thread. This rule is mandated by the EJB Specification 1.1. Also, note that the EJB Specification discourages the direct use of threads by EJB programmers.

Note: If a passivated stateful bean is not removed due to application or system errors, its passivated state takes up disk space in the location specified by the `<persistence-store-directory-root>` element. This passivated state remains on disk until the temporary files containing that state are deleted by the System Administrator. These files can be identified by the syntax of their names, which include the following information:

- The server name, which includes the server group name and the group ID
- The server generation ID
- The bean name
- A string of several digits

The server name and bean name components of the filename are the most readily identifiable. To manage the number of unused bean state files that can potentially accumulate over time, System Administrators may choose to create scripts that delete those files whenever the WebLogic Enterprise system is started or shut down.

Controlling the Passivation of Beans -- the EJB Cache

In WebLogic Enterprise 5.0, the EJB container cached stateful beans only for the duration of a transaction. If such beans were not involved in a transaction, the container passivated them after each method invocation on the bean.

In versions of WebLogic Enterprise after 5.0, the EJB container has the ability to cache beans across method invocations as well as across transactions. This significantly reduces the frequency of beans being passivated, thereby providing the performance improvement. EJB caching is enabled by default for stateful beans.

When Stateful Beans Are Passivated

A cached stateful bean is normally only stored (passivated) if it is unused for a period of time. You can configure an optional cache flush time, if desired. A bean may be considered unused if other beans are being used more frequently and the bean is the least recently used bean. The bean may also be passivated if the cache flush time occurs and the bean is not presently active within a method call or a transaction.

Passivation After Creation

In the WebLogic Enterprise system, the EJB container still passivates stateful beans immediately after the `ejbCreate` method is called. This behavior may change in future releases, however.

How to Set Up EJB Caching

EJB caching is enabled by default for stateful session beans and entity beans in the EJB container of the WebLogic Enterprise system, and can be set up using the following mechanisms:

- Via the WebLogic EJB extensions to the deployment descriptor DTD. Using this mechanism allows you to disable caching for individual beans.

- Specifying the JavaServer parameters `MAXEJBCACHE` and `EJBCACHE_FLUSH`, in which you can establish two settings for bean caching:
 - Maximum number of beans that can be cached at any one time
 - A time interval specifying when the bean cache is to be flushed by the system

For more information about setting up the bean cache, see the section “Starting JavaServer” in the topic [Creating a Configuration File](#). For information about scaling and tuning the bean cache, see [Scaling, Distributing, and Tuning Applications](#).

EJBs as Client Applications

A bean may invoke the methods of another bean. When a bean behaves as a client application, the client rules still apply: the bean must obtain the reference to the other bean from that bean’s home interface (factory), and references to the home interface must be obtained using JNDI.

The main differences with the client environment are the following:

- When the bean creates the `InitialContext` object, there is no authentication or connection setup because WebLogic Enterprise Java servers run within the trusted server base.
- The WebLogic Enterprise EJB container does not support reentrancy, and rejects loopback calls (a bean calling another bean that then calls the first bean) by throwing an exception (`java.rmi.RemoteException`) to the client application.

Note: The WebLogic Enterprise EJB container does not propagate the security and transaction context on callbacks to a J2EE client.

Security, Transactions, and JDBC Connections

For additional EJB design considerations, see the following in the WebLogic Enterprise online documentation:

- [Using Security](#)
- [Using Transactions](#)
- [Using the JDBC Drivers](#)

EJBs and Persistence

This section provides a general discussion about EJBs and persistence, and shows sample fragments of EJB implementation code and deployment descriptors to illustrate the following:

- Container-managed Entity Beans
- Bean-managed Entity Beans
- Stateful Session Beans
- Stateless Session Beans

The topics described in this section use code from the EJB sample applications that are installed with the WebLogic Enterprise software.

Note: The code fragments shown in this section are taken from the EJB samples provided with the WebLogic Enterprise software. These samples include tracing code, which is turned on by the `VERBOSE` flag, that helps show what is happening when the samples are executed. These tracing statements are not required by the EJB Specification 1.1; they are present for instructional purposes only.

Development Considerations for EJBs and Persistence

Persistence refers to a bean's state information, which may be contained in durable storage when the bean is not active. When the bean is activated, this state is read in from durable storage. As a Bean Provider, you basically have two choices for what kind of broad mechanism you want to use for handling a bean's persistence: either directly in the bean's logic, or by delegating to the EJB container the tasks of handling the bean's persistence.

A bean that delegates to the container all the logic for handling its persistent data has what is referred to as **container-managed persistence**. A bean that contains its own logic for handling its persistence data has what is referred to as **bean-managed persistence**. The choice you make for a bean must be specified in the bean's deployment descriptor.

Container-managed Entity Beans

If your entity bean uses container-managed persistence, you need to do the following:

- In the bean's standard deployment descriptor, define the elements described in the section "Required Deployment Descriptor Elements for Container-managed Beans" on page 8-12.
- In the bean's implementation, declare as public the class variables whose persistence is container-managed, as shown in the code fragment in the section "Declaring Container-managed Fields as Public Variables" on page 8-13.

Note: The EJB Specification 1.1 requires that class variables whose persistence is container-managed have public access.

The subsections that follow also provide code fragments that show the use of the `ejbCreate`, `ejbStore`, and `ejbRemove` methods in such beans.

Required Deployment Descriptor Elements for Container-managed Beans

The required deployment descriptor elements for container-managed beans are listed and described in Table 8-3.

Table 8-3 Required Deployment Descriptor Elements for Container-managed Beans

Deployment Descriptor Element	Description
<code>cmp-fields</code>	This element specifies the container-managed fields. This is a standard property that lists the public nontransient instance variables that the EJB expects will be made automatically persistent. Even if there are no managed fields, the bean's object reference and the primary key are remembered by the EJB container.
<code>persistence-type</code>	If the EJB is an entity bean, this element declares whether the persistence management is performed by the container or by the bean. For container-managed persistence, this element should specify <code>container</code> .

Table 8-3 Required Deployment Descriptor Elements for Container-managed Beans (Continued)

Deployment Descriptor Element	Description
<code>persistence-store-jdbc</code>	This element specifies that the type of persistent store is <code>jdbc</code> . (Note that in WebLogic Enterprise, file-based persistence for entity beans is not supported.)
<code>primary-key-class</code>	The deployment descriptor for any entity bean must set this element, which identifies the primary key class for the bean.

Declaring Container-managed Fields as Public Variables

The following code fragment shows a container-managed EJB declaring its container-managed fields as public class variables. These variables need to be public so that the container can manage them.

```
// public container managed variables
public          String      accountId; // also the primary Key
public          double      balance;
public          String      type; // "Checking"
```

The ejbCreate Method

A container-managed entity bean needs to implement the `ejbCreate` method. Note that the `accountId` and `initialBalance` parameters in this method are managed by the container. The following code fragment shows a container-managed bean setting the values of the public class variables shown in the code fragment in the section “Declaring Container-managed Fields as Public Variables” on page 8-13.

```
public void ejbCreate(String accountId, double initialBalance,
                      String type) {
    if (VERBOSE)
        System.out.println("ejbCreate( id = " + id() +
                           ", initial balance = $" + initialBalance +
                           ", type: " + type + ")");
    this.accountId = accountId;
    this.balance   = initialBalance;
    this.type      = type;
}
```

The ejbStore Method

The EJB Specification 1.1 states that implementing the `ejbStore` method in a container-managed entity bean is required, even if the method does not provide any specific functionality. One advantage to having this method in your container-managed bean is to provide a tracing capability for debugging purposes, as in the following example:

```
public void ejbStore() {
    if (VERBOSE)
        System.out.println("ejbStore (" + id() + ")");
}
```

The ejbRemove Method

As with the `ejbStore` method in the preceding section, the `ejbRemove` method is not a functional component of a container-managed bean implementation; however, the EJB Specification 1.1 requires this method to be present, as in the following example:

```
public void ejbRemove() throws RemoveException {
    if (VERBOSE)
        System.out.println("ejbRemove (" + id() + ")");
}
```

Bean-managed Entity Beans

If you are implementing a bean with bean-managed persistence, you need to do the following:

- Declare the bean's persistence type in the persistence-type element in the standard deployment descriptor.
- Implement code in the bean that accesses the JDBC connection pool.
- Implement the `ejbCreate` and `ejbStore` methods to create a database entity that contains the bean's persistent data, and store that data.

The code fragments provided in this section illustrate performing these tasks, as well as using the `ejbRemove` method to remove a bean's persistent data from a database.

Accessing the JDBC Pool

The following code example shows a bean-managed entity bean using static initialization to establish access to the JDBC pool, which is defined in the EJB application's UBBCONFIG file:

```
static {
    try{
        Context ctx = new InitialContext();
        pool = (DataSource)ctx.lookup("jdbc/pool1");
    } catch(Exception e) {
        System.out.println("problem with datasource.");
    }
}
```

The ejbCreate Method

A bean-managed entity bean uses the `ejbCreate` method to create the bean and update the table in the database that contains the entity bean's value. Listing 8-1 shows creating a row in the table in that database, using a JDBC connection from the pool.

Listing 8-1 Entity Bean Creating a Row in a Database

```
public AccountPK ejbCreate(String account_id, double initial_balance)
    throws CreateException,
{
    if (VERBOSE) {
        System.out.println("AccountBean.ejbCreate( id = " +
                           System.identityHashCode(this) + ", PK = " +
                           account_id + ", " + "initial balance = $" +
                           initial_balance + ")");
    }
    AccountId = account_id;
    Balance = initial_balance;

    Connection connection = null;
    PreparedStatement prep_stmt = null;

    try {
        connection = pool.getConnection();
        prep_stmt = connection.prepareStatement("insert into ejbAccounts "+
                                              "(id, bal) values (?, ?)");
        prep_stmt.setString(1, AccountId);
    }
```

```
prep_stmt.setDouble(2, Balance);
if (prep_stmt.executeUpdate() != 1) {
    throw new CreateException ("JDBC did not create any row");
}
AccountPK primary_key = new AccountPK();
primary_key.AccountId = AccountId;
return primary_key;
} catch (CreateException ce) {
    throw ce;
} catch (SQLException sqe) {
    throw new CreateException (sqe.getMessage());
} finally {
    try {
        prep_stmt.close();
        connection.close();
    } catch (Exception e) {
    }
} // end of finally
} // end of ejbCreate(..)
```

Updating the Database

The following code fragment shows updating the database with the values. Since this bean uses bean-managed persistence, updating the database is done manually. Whereas the code in the previous example created the database rows, the code in Listing 8-2 specifies the values in those rows.

Listing 8-2 Entity Bean Updating the Database

```
ejbStore()
public void ejbStore() throws EJBException {

    if (VERBOSE) {
        System.out.println("ejbStore (" + id() + ")");
    }
    Connection connection = null;
    PreparedStatement prep_stmt = null;
    try {
        connection = pool.getConnection();
        prep_stmt = connection.prepareStatement("update ejbAccounts set bal = "+
                                             "? where id = ?");
        prep_stmt.setDouble(1, Balance);
        prep_stmt.setString(2, AccountId);
        int i = prep_stmt.executeUpdate();
        if (i == 0) {
```

```
        throw new RemoteException ("ejbStore: AccountBean (" + AccountId +
                                  ") not updated");
    }
} catch (RemoteException re) {
    throw re;
} catch (SQLException sqe) {
    throw new EJBException (sqe.getMessage());
} finally {
    try {
        prep_stmt.close();
        connection.close();
    } catch (Exception e) {
        throw new EJBException (e.getMessage());
    }
} // end of finally
} // end of ejbStore()
```

Removing Values From the Database

Listing 8-3 shows using the `ejbRemove` method to remove rows from the database that were created and set in the preceding code examples.

Listing 8-3 Removing Values From a Database

```
public void ejbRemove()
    throws RemoveException,
{
    if (VERBOSE) {
        System.out.println("ejbRemove (" + id() + ")");
    }
    // we need to get the primary key from the context because
    // it is possible to do a remove right after a find, and
    // ejbLoad may not have been called.

    Connection connection = null;
    PreparedStatement prep_stmt = null;
    try {
        connection = getConnection();
        AccountPK pk = (AccountPK) ctx.getPrimaryKey();
        prep_stmt = connection.prepareStatement("delete from ejbAccounts where "+
                                              "id = ?");
        prep_stmt.setString(1, pk.AccountId);
        int i = prep_stmt.executeUpdate();
        if (i == 0) {
            throw new EJBException ("AccountBean (" +

```

```
        + pk.AccountId + " not found");
    }
} catch (RemoteException re) {
    throw re;
} catch (SQLException sqe) {
    throw new RemoteException (sqe.getMessage());
} finally {
    try {
        prep_stmt.close();
        connection.close();
    } catch (Exception e) {
        throw new EJBException (e.getMessage());
    }
} // end of finally
} // end of ejbRemove()
```

Stateful Session Beans

This section shows the following examples of stateful session beans:

- Required standard deployment descriptor elements
- Code fragments showing two stateful session beans: one in which the client keeps track of the bean's state, and one in which the bean keeps track of its state

The code examples shown here are from the EJB Samples directory, which is available with the WebLogic Enterprise software.

Example Deployment Descriptor

A deployment descriptor for a stateful session bean can optionally define the `persistentDirectoryRoot` element. The default file is `/pstore/bean_name.dat`, where the directory `pstore` represents the directory from which the WebLogic Enterprise application was started, and `bean_name` is the fully qualified name of the EJB with underscores (`_`) replacing the periods (`.`) in the name.

If the `persistentStoreType` element is defined as `jdbc`, the container looks for additional values to determine the appropriate values for the JDBC connection. Note that if the bean's persistence is stored in a database via a JDBC connection, the System Administrator needs to add this information to the `UBBCONFIG` file as well. For more information, see [Using the JDBC Drivers](#) in the WebLogic Enterprise online documentation.

The following deployment descriptor fragment shows the location of the persistent store for a stateful session bean:

```
<persistence-store-descriptor>
    <persistence-store-file>
        <persistence-store-directory-root>
            c:\mystore
        </persistence-store-directory-root>
    </persistence-store-file>
</persistence-store-descriptor>
```

Client Application Maintaining a Bean's State Information

The following code example shows an EJB client application keeping track of a bean's state information. In stateful session beans, you need to provide a one-to-one mapping between the client and the bean in the server, represented by a key. This key provides the map between the bean's instance and the client, because the bean instance cannot be shared with other clients.

Listing 8-4 shows the client application code creating the stateful session bean using the primary class key.

Listing 8-4 Client Application Using the Primary Class Key

```
// Give this trader a name
Trader trader = brokerage.create("Terry");
System.out.println("Creating trader " + trader.getTraderName() + "\n");

String stockName;
int    numberOfShares;

for (int i = 1 ; i <= 5; i++) {
    System.out.println("Start of Transaction " + i + " for " + customerName);

    // Buying
    stockName      = "WEBL";
    numberOfShares = 100 * i;
    System.out.println("Buying " + numberOfShares + " of " + stockName);
    TradeResult tr = trader.buy(customerName, stockName, numberOfShares);
    System.out.println("...Bought " + tr.numberTraded + " at $" +
                      tr.priceSoldAt);

    // Selling
    stockName      = "INTL";
    numberOfShares = 100 * (i+1);
```

```
System.out.println("Selling " + numberOfShares + " of " + stockName);
tr = trader.sell(customerName, stockName, numberOfShares);
System.out.println("...Sold " + tr.numberTraded + " at $" +
tr.priceSoldAt);

// Get change in Cash Account from EJBean
System.out.println("Change in Cash Account: $" + trader.getBalance());

System.out.println("End of Transaction " + i + "\n");
}
System.out.println("Change in Cash Account: $" + trader.getBalance() + "\n");
System.out.println("Removing trader " + trader.getTraderName());
trader.remove();
}
catch (ProcessingErrorException pe) {
    System.out.println("Processing Error: " + pe);
    pe.printStackTrace();
}
catch (Exception e) {
    System.out.println("::::::::::::: Error :::::::::::::::");
    e.printStackTrace();
}
```

Bean Keeping Track of Its Own State

Listing 8-5 shows a stateful session bean keeping track of its state, and its mapping to a specific client. For example, the balance is kept on the EJB rather than on the client.

Listing 8-5 EJB State Management

```
// The reason the following attribute is public is to test
// passivation into a persistent store, because the deployment descriptor
// says it should be a stateful session bean.
// This and the ejbCreate method in this file are the differences
// between the examples in the stateful and stateless directories.
public String traderName;
public double tradingBalance;

// -----
.

.

public TradeResult buy(String customerName, String stockSymbol,
int shares)
```

```

    throws ProcessingErrorException
{
    if (VERBOSE && shares >= 0) {
        System.out.println("buy (" + customerName + ", " +
                           stockSymbol + ", " +
                           shares + ")");
    }
    try {
        int tradeLimit = getTradeLimit();
        if (shares > tradeLimit)          // limit for buying
            shares = tradeLimit;
        else if (shares < -tradeLimit) // limit for selling
            shares = -tradeLimit;

        double price = getStockPrice(stockSymbol);
        tradingBalance = tradingBalance - (shares * price); // subtract purchases
from cash account

        if (shares < 0)
            shares = -shares;
        return new TradeResult(shares, price);
    }
    catch (Exception e) {
        throw new ProcessingErrorException("Trader error: " + e);
    }
}

```

Stateless Session Beans

This section provides the following two code examples:

- An EJB client application keeping track of a stateless session bean's state
- A stateless bean that keeps track of its own state data

Client Maintaining Bean's State

Listing 8-6 shows a client application keeping track of the `cashBalance` variable, which is manipulated by the stateless bean. This example also shows the client invoking the `ejbCreate` method without any arguments and without any specific data.

Listing 8-6 Client State Management

```
try {
    String  customerName  = "Erin";      // Default name for the customer
    Context  ctx          = null;        // To hold JNDI context
    Object  objref        = null;        // to hold object reference

    String  stockName     = null;        // Name of a stock
    int     numberOfShares = 0;          // No. of shares to trade
    double  cashBalance   = 0.0;        // To hold balance between sessions

    TraderHome brokerage   = null;        // To hold home interface
    Trader    trader       = null;        // To hold trader object
    TradeResult tradeResult = null;      // To hold results from a trade

    // Create a new initial context based on the url, user, and password
    ctx = newInitialContext();

    if ( ctx == null ){
        System.out.println("Initial context is null");
        exit(-1);
    }
    // do a JNDI lookup for the EJB;defined in the deployment descriptor
    objref = ctx.lookup("statelessSession.TraderHome");
    printTrace("Looked up home:");

    /* Create a trader object, who'll later help us execute trades
     * The lookup has resulted in an Object.  We know
     * this object is actually a reference of type TraderHome,
     * so the reference is narrowed and cast to that type:
     */
    brokerage = (TraderHome) PortableRemoteObject.narrow(objref,
                                                          TraderHome.class);
    printTrace("Narrowed home.");

    /* Create the EJB on the WLE server.
     * Unlike the statefulSession example,we don't give this trader a key
     */
    printTrace("Creating trader.");
    trader = brokerage.create();

    /* Unlike the statefulSession example,
     * we have to keep track of the balance over the
     * life of our use of the session bean
     */

    for (int i = 1 ; i <= maxTransaction; i++) {
        System.out.println("Start of Transaction " + i + " for " +
```

```
    customerName);  
  
    /* Buying  
     * Stock symbol must be found in the deployment descriptor's environment  
     * properties section. TraderBean EJB will check the validity of the  
     * symbol, and its price using JNDI lookup on the environment  
     * properties.  
     */  
  
    stockName      = "BEAS";  
    numberOfShares = 100 * i;  
    System.out.println("Buying " + numberOfShares + " of " + stockName);  
  
    // buy() is executed on the TraderBean EJB in the WLE Server  
    tradeResult = trader.buy(customerName, stockName, numberOfShares);  
    System.out.println("...Bought " + tradeResult.numberTraded + " at $" +  
                      tradeResult.priceSoldAt);  
  
    // Keep track of the change in the Cash Account  
    cashBalance = cashBalance - (tradeResult.numberTraded *  
                                 tradeResult.priceSoldAt);  
  
    // Selling  
    stockName      = "INTL";  
    numberOfShares = 100 * (i+1);  
    System.out.println("Selling " + numberOfShares + " of " + stockName);  
  
    // sell() is executed on the TraderBean EJB in the WLE Server  
    tradeResult = trader.sell(customerName, stockName, numberOfShares);  
    System.out.println("...Sold " + tradeResult.numberTraded + " at $" +  
                      tradeResult.priceSoldAt);  
  
    // Keep track of the change in the Cash Account  
    cashBalance = cashBalance + (tradeResult.numberTraded *  
                                 tradeResult.priceSoldAt);  
  
    // Print change in Cash Account  
    System.out.println("Change in Cash Account: $" + cashBalance);  
    System.out.println("End of Transaction " + i + "\n");  
}  
System.out.println("Change in Cash Account: $" + cashBalance + "\n");  
System.out.println("Removing trader");  
  
// Remove TraderBean EJB from the WLE server.  
trader.remove();  
}  
catch (ProcessingErrorException pe) {  
    System.out.println("Processing Error: " + pe);  
    pe.printStackTrace();
```

```
    }
    catch (Exception e) {
        System.out.println(":::::::::::::::::: Error ::::::::::::::::::::");
        e.printStackTrace();
    }
}
```

Stateless Bean Tracking Its Own State

Listing 8-7 shows a business method from the `TraderBean` example, available in the EJB Samples directory provided with the WebLogic Enterprise software. In this example, the bean does not preserve any state. The bean's `buy` method performs simple calculations on data provided by the client application.

Listing 8-7 Stateless Bean State Management

`getStockPrice()` and `getTradeLimit()` methods use DD environment properties to access constant values using JNDI `lookup()` - prevents hardcoding data.

```
public TradeResult buy(String customerName, String stockSymbol,
                      int shares)
    throws ProcessingErrorException
{
    if (shares >= 0) {
        printTrace("buy (" + customerName + ", " +
                  stockSymbol + ", " +
                  shares + ")");
    }
    try {
        int tradeLimit = getTradeLimit();
        if (shares > tradeLimit)
            shares = tradeLimit;
        else if (shares < -tradeLimit) // limit for selling
            shares = -tradeLimit;

        double price = getStockPrice(stockSymbol);
        printTrace("Executing buy...");
        if (shares < 0)
            shares = -shares;
        return new TradeResult(shares, price);
    }
    catch (Exception e) {
        throw new ProcessingErrorException("Trader error: " + e);
    }
}
```

Index

A

activation 8-7
activation policies
 defining in Implementation
 Configuration file 4-16
 defining in Server Description file 4-16
 Simpapp sample application 4-16
 Simple interface 4-17
 supported 4-16
ActiveX application builder
 description 2-6
AdminAPI
 description 2-5
administration commands
 tmadmin command 2-4
 tmboot command 2-4
 tmconfig command 2-4
 tmloadcf command 2-4
 tmshutdown command 2-4
 tmunloadcf command 2-4
Administration console
 description 2-4
administration tools
 AdminAPI 2-5
 administration commands 2-4
 Administration console 2-4
app 4-26
Application Assembler 3-5
Authenticates 2-18
authentication
 client application 2-18

levels 5-2

B

bean-managed persistence 3-5
beans
 See EJBs
Bootstrap object
 description 2-11
 illustrated 2-11
 Simpapp sample application 4-21
building
 C++ client applications 4-30
 buildobjclient command 2-3
 C++ server applications
 buildobjserver command 2-3
 genifc command 2-3
 Java client applications 4-30
 Java server applications
 buildjavaserver command 2-3
buildjavaserver command
 building Java server applications 2-3
 description 2-3
 format 4-29
 in the Simpapp sample application 4-29
buildobjclient command
 building C++ client applications 2-3
 description 2-3
 format 4-30
 in the Simpapp sample application 4-30
buildobjserver command

- building C++ server applications 2-3
- description 2-3
- format 4-29
- in the Simpapp sample application 4-29
- buildXAJS** command
 - building an XA resource manager 2-3
 - description 2-3

- C**
- client 4-30
- client applications
 - authenticating into the WLE domain 2-18
 - EJB 8-2
 - initialization process 2-18
 - invoking objects 2-21
 - using transactions 6-4
 - writing
 - Security sample application 6-11
 - Simpapp sample application 4-21
 - Transactions sample application 6-11
- client stubs
 - generating 4-6
 - in Simpapp sample application 4-6
- code example
 - C++ client application for Simpapp sample application 4-22
 - C++ implementation of the Simple interface 4-9
 - C++ Server object 4-12
 - C++ server object that supports transactions 6-13
 - configuration file for Simpapp sample application 4-26
 - Java client application for the Simpapp sample application 4-23
 - Java implementation of SimpleFactory interface 4-11
 - Java implementation of the Simple
- Interface 4-10
- Java Server object 4-14
- OMG IDL for Simpapp sample application 4-6
- OMG IDL for Transactions sample application 6-8
- security in C++ client applications 5-7
- security in Java client applications 5-8
- transactions in C++ client application 6-12
- UBBCONFIG file for Transactions sample application 6-15
- compiling
 - C++ client applications 4-30
 - C++ server applications 4-29
 - Java client applications 4-30
 - Java server applications 4-29
- container-managed persistence 3-5
- CORBAservices Object Transaction Service
 - using in WLE applications 6-2
- create_servant method 2-19
- customer support contact information ix

D

- Deployer 3-5
- deployment, dynamic
 - See hot redeployment
- development commands
 - buildjavaserver 2-3
 - buildobjclient command 2-3
 - buildobjserver command 2-3
 - buildXAJS command 2-3
 - genifc command 2-3
 - idl2ir command 2-3
 - ir2idl command 2-3
 - irdel command 2-3
- development process
 - activation policies 4-16
 - client applications
 - Security sample application 5-7

Simpapp sample application 4-21
Transactions sample application 6-11
defining object activation policies 4-16
illustrated 4-3
Implementation Configuration file 4-16
OMG IDL
 Simpapp sample application 4-5
 Transactions sample application 6-8
Security sample application 5-6
server applications
 Simpapp sample application 4-8
 Transactions sample application 6-13
Server Description file 4-16
Simpapp sample application 4-4
steps for creating WLE applications 4-2
Transactions sample application 6-7
WLE applications 4-2
writing a configuration file 4-25
writing server application code 4-8
writing the client application code 4-21
writing the OMG IDL 4-5
documentation, where to find it viii
dynamic deployment
 See hot redeployment

E

ejbCreate method 8-5
ejbPostCreate method 8-5
ejbRemove method 8-5
EJBs
 and persistence 3-5
 as clients 8-10
 Container Provider 3-5
 Deployer 3-5
 design considerations 8-5
 designing for client applications 8-2
 entity 3-3, 8-6
 initializing in server 8-7

provider 3-5
Server Provider 3-5
session 8-6
stateful session 3-3
stateless session 3-3
Enterprise Bean Provider 3-5
Enterprise JavaBeans
 see EJBs
entity beans 3-3, 8-6
EntityBean interface 8-5
environmental objects
 and client initialization 2-18
 description 2-8

F

factories
 finding 2-19
 registering 2-19
FactoryFinder object
 description 2-8
 example use of 2-19

G

genicf command
 creating a ICF file 2-3
 description 2-3

H

Home interface 8-2
host
 port 8-2
hot redeployment 7-37
 and the module initializer object 7-13

I

IDL
 See Interface Definition Language 2-2
idl command 2-2

description 2-2
files created by 4-7
generating client stubs 4-7
generating skeletons 4-7
IDL compiler
 idl command 2-2
 m3idltojava command 2-2
 supported 2-2
idl2ir command
 description 2-3
idltojava compiler
 differences from Sun version 2-2
IIOP Listener/Handler
 description 2-12
Implementation Configuration file
 defining activation policies 4-16
 defining transaction policies 6-10
initialize method
 on Server object 8-7
 summary 2-16, 2-17
initializing EJB applications 8-7
Interface Repository
 creating 2-3
 deleting objects from 2-3
 displaying the contents 2-3
 idl2ir command 2-3
 ir2idl command 2-3
 irdel command 2-3
 loading interface definitions into 2-3
InterfaceRepository object
 description 2-9
interfaces
 writing methods to implement
 operations 4-9
ir2idl command
 description 2-3
irdel command
 description 2-3

J
Java client applications
 required files 4-30
Java Transaction Service
 using in WLE applications 6-2

K
key
 primary 3-3

M
m3idltojava command 2-2
 description 2-2
 files created by 4-8
 generating client stubs 4-7
 generating skeletons 4-7
Management Information Base
 see MIB 1-5
managing
 WLE applications
 tmadmin command 2-4
 tmboot command 2-4
 tmconfig command 2-4
 tmloadcf command 2-4
 tmshutdown command 2-4
 tmunloadcf command 2-4
method implementations
 C++ 4-9
 Java 4-10
 writing 4-9
MIB
 for WLE applications 1-5

O
object handle 8-2
Object Life Cycle service
 description 2-8
object request broker

- see ORB 2-13
- object services
 - Interface Repository 2-9
 - Object Life Cycle service 2-8
 - Security service 2-8
 - Transaction service 2-8
- objects
 - invoking 2-21
- OMG IDL
 - compiling 4-6
 - generating client stubs 4-6
 - generating skeletons 4-6
 - Simple interface 4-5, 4-6
 - SimpleFactory interface 4-5, 4-6
 - Transactions sample application 6-8
- ORB
 - description 2-13
 - illustrated 2-13

P

- passivation 8-7
- persistence
 - bean-managed 3-5
 - container-managed 3-5
- POA
 - description 2-14
 - interaction with TP Framework 2-14
- Portable Object Adapter
 - see POA 2-14
- primary key 3-3
- PrincipalAuthenticator object
 - using in client applications 5-4
- printing product documentation ix
- programming tools 2-2

R

- register_factory method
 - example of 2-19
- related information ix

- release method
 - on Server object 8-7
- Remote interface 8-2
- RemoteException exception 8-10
- resolve_initial_references method 2-18

S

- security authentication 8-2
- security credentials 8-2
- security principal 8-2
- Security sample application
 - defining security level 5-6
 - description 5-4
 - development process 5-6
 - illustrated 5-4
 - location of files 5-5
 - PrincipalAuthenticator object 5-4
 - SecurityCurrent object 5-4
 - using the PrincipalAuthenticator object 5-7
 - using the SecurityCurrent object 5-7
 - writing the client application 5-7
- Security service
 - description 2-8
 - functional description 5-2
- SecurityCurrent object
 - description 2-8
 - using in client applications 5-4
- server applications
 - defining object activation policies 4-16
 - Implementation Configuration file 4-16
 - Server Description file 4-16
 - writing
 - Simpapp sample application 4-8
 - Transactions sample application 6-13
 - writing method implementations 4-9
 - writing the Server object 4-12
- Server Description file
 - defining activation policies 4-16

defining transaction policies 6-11

Server object 6-13, 8-7

- description 2-16
- Transactions sample application 6-13
- writing 4-12

session beans 3-3, 8-6

SessionBean interface 8-5

SessionSynchronization interface 8-5

Simpapp sample application

- compiling
 - C++ client application 4-30
 - C++ server application 4-29
 - Java client application 4-30
- compiling Java server application 4-29
- configuration file 4-25, 7-29
- description 4-4
- file location 4-5
- illustrated 4-4
- interfaces defined for 4-5
- OMG IDL 4-5
 - using the Bootstrap object 4-21
 - using the buildjavaserver command 4-29
 - using the buildobjserver command 4-29
 - writing the client application code 4-21

Simple interface

- activation policy 4-17
- OMG IDL 4-5

Simple Network Management Protocol

- see SNMP 1-5

SimpleFactory interface

- OMG IDL 4-5

skeletons

- generating 4-6
- in Simpapp sample application 4-6

SNMP

- in the WLE product 1-5
- stateful 8-6
- stateful session beans 3-3
- stateless 8-6
- stateless session beans 3-3
- support

technical x

supporting databases 6-13

System Administrator 3-5

T

T_MODULE TMIB class 7-37

TLOGDEVICE parameter 6-15

tmadmin command

- description 2-4

tmboot command

- description 2-4

tmconfig command

- description 2-4

tmloadcf command

- creating a configuraiton file 4-28, 7-33
- description 2-4

tmshutdown command

- description 2-4

tmunloadcf command

- description 2-4

Tobj_Bootstrap 2-18

TP Framework

- description 2-14
- illustrated 2-15

transaction policies

- defined 6-10

Transaction server application

- writing the server application 6-13

Transaction service

- description 2-8, 6-1
- features 6-2

TransactionCurrent object

- description 2-8

transactions

- functional overview 6-3
- illustrated 6-3
- in client applications 6-4
- OMG IDL 6-4

Transactions sample application

- description 6-4

- file location 6-7
- illustrated 6-5
- OMG IDL 6-8
- starting server application 6-13
- transaction policies 6-11
- UBBCONFIG file 6-14
- writing client applications 6-11
- writing server applications 6-13
- Transactions sample application
 - development process 6-7
- UXCONFIG file
 - description 4-28, 7-33

U

- UBBCONFIG file
 - adding transactions 6-14
 - description 4-28, 7-33
 - sections in 4-25
 - setting the security level 5-6
- user exceptions
 - Transactions sample application 6-5

UserTransaction object

- description 2-8

W

- WLE applications
 - defining security levels 5-6
 - how they work 2-16
 - managing
 - tmadmin command 2-4
 - tmboot command 2-4
 - tmconfig command 2-4
 - tmloadcf command 2-4
 - tmshutdown command 2-4
 - tmunloadcf command 2-4
 - using CORBAservices Object
 - Transaction Service 6-2
 - using Java Transaction Service 6-2
- WLE components

