

Oracle® Fusion Middleware

Developing JNDI Applications for Oracle WebLogic Server

12c (12.1.2)

E28142-03

March 2015

This document explains how to use the JNDI API to provide a unified interface to multiple naming and directory services in the enterprise.

Copyright © 2007, 2015, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	v
Documentation Accessibility	v
Conventions	v
1 Developing JNDI Applications for Oracle WebLogic Server	
1.1 Contents.....	1-1
1.2 Introduction and Roadmap	1-1
1.2.1 Guide to this Document.....	1-1
1.2.2 Related Documentation	1-2
1.2.3 Examples for the Web Application Developer.....	1-2
1.2.3.1 Avitek Medical Records Application (MedRec)	1-2
1.2.3.2 JNDI Examples in the WebLogic Server Distribution.....	1-2
1.2.4 New and Changed Features in This Release	1-2
1.3 Understanding WebLogic JNDI.....	1-2
1.3.1 What is JNDI?	1-3
1.3.2 WebLogic Server JNDI.....	1-3
1.4 WebLogic JNDI	1-3
1.4.1 Using WebLogic JNDI to Connect a Java Client to a Single Server.....	1-4
1.4.2 Setting Up JNDI Environment Properties for the Initial Context.....	1-4
1.4.2.1 Creating a Context Using a Hashtable	1-6
1.4.2.2 Creating a Context Using a WebLogic Environment Object	1-6
1.4.2.3 Creating a Context from a Server-Side Object.....	1-7
1.4.2.4 Associating a WebLogic User with a Security Context.....	1-7
1.4.2.5 JNDI Contexts and Threads	1-7
1.4.2.5.1 How to Avoid Potential JNDI Context Problems.....	1-7
1.4.2.5.2 Last Used	1-7
1.4.3 Using the Context to Look Up a Named Object	1-8
1.4.4 Using a Named Object to Get an Object Reference.....	1-8
1.4.5 Closing the Context	1-8
1.4.6 Using WebLogic JNDI in a Clustered Environment	1-9
1.4.6.1 Using the Relationship of RMI and JNDI to Enable WebLogic Clusters	1-9
1.4.6.2 Making Custom Objects Available to a WebLogic Server Cluster.....	1-9
1.4.6.3 Data Caching Design Pattern.....	1-10
1.4.6.4 Exactly-Once-Per-Cluster Design Pattern.....	1-11
1.4.6.5 Using WebLogic JNDI from a Client in a Clustered Environment.....	1-11

1.4.7	Using JNDI from Within Java EE Components.....	1-13
1.4.8	Setting Up Foreign JNDI.....	1-14

Preface

This preface describes the document accessibility features and conventions used in this guide—*Understanding the WebLogic Scripting Tool*.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Conventions

The following text conventions are used in this document:

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

Developing JNDI Applications for Oracle WebLogic Server

This document explains how to set up WebLogic JNDI. It is intended for programmers who are developing WebLogic Server applications and want to use the JNDI feature.

This document is written for application developers who want to design, develop, configure, and manage applications using the Java Platform, Enterprise Edition (Java EE) and want to use the JNDI API to provide a unified interface to multiple naming and directory services in their enterprise. It is assumed that readers know JNDI and the Java programming language.

1.1 Contents

This document contains the following topics:

- [Section 1.2, "Introduction and Roadmap"](#)
- [Section 1.3, "Understanding WebLogic JNDI"](#)
- [Section 1.4, "WebLogic JNDI"](#)

1.2 Introduction and Roadmap

The following sections describe the contents and organization of this guide—*Developing JNDI Applications for Oracle WebLogic Server*.

- [Section 1.2.1, "Guide to this Document"](#)
- [Section 1.2.2, "Related Documentation"](#)
- [Section 1.2.3, "Examples for the Web Application Developer"](#)
- [Section 1.2.4, "New and Changed Features in This Release"](#)

1.2.1 Guide to this Document

- [Section 1.2, "Introduction and Roadmap,"](#) describes the scope and organization of this guide.
- [Section 1.3, "Understanding WebLogic JNDI,"](#) provides an overview of the Java Naming and Directory Interface (JNDI) implementation in WebLogic Server.
- [Section 1.4, "WebLogic JNDI,"](#) describes programming with WebLogic JNDI.

1.2.2 Related Documentation

For additional information on JNDI, see the following documents:

- "JNDI Subsystem Messages" provides a list of JNDI subsystem messages.
- "Communications in a Cluster" in *Administering Clusters for Oracle WebLogic Server* provides information on the cluster-wide JNDI tree.
- *Oracle WebLogic Server Administration Console Online Help* contains sections that describe how to add or modify security roles and policies on a JNDI Binding node, Root Content node, or Context node.

1.2.3 Examples for the Web Application Developer

In addition to this document, Oracle provides examples for software developers within the context of the Avitek Medical Records Application (MedRec) sample application, as well as JNDI code examples.

1.2.3.1 Avitek Medical Records Application (MedRec)

MedRec is an end-to-end sample Java EE application shipped with WebLogic Server that simulates an independent, centralized medical record management system. The MedRec application provides a framework for patients, doctors, and administrators to manage patient data using a variety of different clients.

MedRec demonstrates WebLogic Server and Java EE features, and highlights Oracle-recommended best practices. MedRec is optionally installed with the WebLogic Server installation. You can start MedRec from the `ORACLE_HOME\user_projects\domains\medrec` directory, where `ORACLE_HOME` is the directory you specified as Oracle Home when you installed Oracle WebLogic Server. For more information, see "Sample Applications and Code Examples" in *Understanding Oracle WebLogic Server*.

1.2.3.2 JNDI Examples in the WebLogic Server Distribution

WebLogic Server optionally installs API code examples in the `EXAMPLES_HOME\wl_server\examples\src\examples` directory, where `EXAMPLES_HOME` represents the directory in which the WebLogic Server code examples are configured. For more information about the WebLogic Server code examples, see "Sample Applications and Code Examples" in *Understanding Oracle WebLogic Server*.

1.2.4 New and Changed Features in This Release

For a comprehensive listing of the new WebLogic Server features introduced in this release, see *What's New in Oracle WebLogic Server*.

1.3 Understanding WebLogic JNDI

The following sections present an overview of the Java Naming and Directory Interface (JNDI) implementation in WebLogic Server including:

- [Section 1.3.1, "What is JNDI?"](#)
- [Section 1.3.2, "WebLogic Server JNDI"](#)

1.3.1 What is JNDI?

Applications use naming services to locate objects in data sources, EJBs, JMS, Mail Sessions, and so on in the network. A naming service associates names with objects and finds objects based on their given names. (The RMI registry is a good example of a naming service.)

JNDI provides a common-denominator interface to many existing naming services, such as LDAP (Lightweight Directory Access Protocol) and DNS (Domain Name System). These naming services maintain a set of bindings, which relate names to objects and provide the ability to look up objects by name. JNDI allows the components in distributed applications to locate each other.

The JNDI API, at

<http://docs.oracle.com/javase/7/docs/technotes/guides/jndi/reference.html>, is defined to be independent of any specific naming or directory service implementation. It supports the use of a number of methods for accessing various new and existing services. This support allows any service-provider implementation to be plugged into the JNDI framework using the standard service provider interface (SPI) conventions.

1.3.2 WebLogic Server JNDI

The WebLogic Server implementation of JNDI supplies methods that:

- Give clients access to the WebLogic Server naming services
- Make objects available in the WebLogic namespace
- Retrieve objects from the WebLogic namespace

Each WebLogic Server cluster is supported by a replicated cluster-wide JNDI tree that provides access to both replicated and pinned RMI and EJB objects. While the JNDI tree representing the cluster appears to the client as a single global tree, the tree containing the cluster-wide services is actually replicated across each WebLogic Server instance in the cluster. For more information, see [Section 1.4.6, "Using WebLogic JNDI in a Clustered Environment."](#)

Other WebLogic services can use the integrated naming service provided by WebLogic Server JNDI. For example, WebLogic RMI can bind and access remote objects by both standard RMI methods and JNDI methods.

In addition to the standard Java interfaces for JNDI, WebLogic Server provides its own implementation, `weblogic.jndi.WLInitialContextFactory`, that uses the standard JNDI interfaces.

You need not instantiate this class directly. Instead, you can use the standard `javax.naming.InitialContext` class and set the appropriate hash table properties, as documented in the section [Section 1.4.2, "Setting Up JNDI Environment Properties for the Initial Context."](#) All interaction is done through the `javax.naming.Context` interface, as described in the JNDI Javadoc.

For instructions on using the WebLogic JNDI API for client connections, see [Section 1.4, "WebLogic JNDI."](#)

1.4 WebLogic JNDI

The following sections describe programming with WebLogic JNDI:

- [Section 1.4.1, "Using WebLogic JNDI to Connect a Java Client to a Single Server"](#)

- [Section 1.4.2, "Setting Up JNDI Environment Properties for the Initial Context"](#)
- [Section 1.4.3, "Using the Context to Look Up a Named Object"](#)
- [Section 1.4.4, "Using a Named Object to Get an Object Reference"](#)
- [Section 1.4.5, "Closing the Context"](#)
- [Section 1.4.6, "Using WebLogic JNDI in a Clustered Environment"](#)
- [Section 1.4.7, "Using JNDI from Within Java EE Components"](#)
- [Section 1.4.8, "Setting Up Foreign JNDI"](#)

1.4.1 Using WebLogic JNDI to Connect a Java Client to a Single Server

The WebLogic Server JNDI Service Provider Interface (SPI) provides an `InitialContext` implementation that allows remote Java clients to connect to WebLogic Server. The client can specify standard JNDI environment properties that identify a particular WebLogic Server deployment and related connection properties for logging in to WebLogic Server.

To interact with WebLogic Server, a Java client must be able to get an object reference for a remote object and invoke operations on the object. To accomplish this, the client application code must perform the following procedure:

1. Set up JNDI environment properties for the `InitialContext`.
2. Establish an `InitialContext` with WebLogic Server.
3. Use the context to look up a named object in the WebLogic Server namespace.
4. Use the named object to get a reference for the remote object and invoke operations on the remote object.
5. Close the context.

The following sections discuss JNDI client operations for connecting to a specific WebLogic Server instance. For information about using JNDI in a WebLogic Server cluster, see [Section 1.4.6.5, "Using WebLogic JNDI from a Client in a Clustered Environment."](#)

Before you can use JNDI to access an object in a WebLogic Server environment, you must load the object into the WebLogic Server JNDI tree.

1.4.2 Setting Up JNDI Environment Properties for the Initial Context

The first task that a Java client application must perform is to create environment properties. The `InitialContext` factory uses various properties to customize the `InitialContext` for a specific environment. You set these properties either by using a hash table or the `set()` method of a WebLogic environment object. These properties, which are specified name-to-value pairs, determine how the `WLInitialContextFactory` creates the context.

The following properties are used to customize the `InitialContext`:

- `Context.PROVIDER_URL`— specifies the URL of the WebLogic Server instance that provides the name service. The default is `t3://localhost:7001`.
- `Context.SECURITY_PRINCIPAL`— specifies the identity of the user (that is, a user defined in a WebLogic Server security realm) for authentication purposes. The property defaults to the `guest` user unless the thread has already been associated with a WebLogic Server user. For more information, see [Section 1.4.2.4, "Associating a WebLogic User with a Security Context."](#)

- `Context.SECURITY_CREDENTIALS`—specifies either the password for the user defined in the `Context.SECURITY_PRINCIPAL` property or an object that implements the `weblogic.security.acl.UserInfo` interface with the `Context.SECURITY_CREDENTIALS` property defined. If you pass a `UserInfo` object in this property, the `Context.PROVIDER_URL` property is ignored. The property defaults to the guest user unless the thread has already been associated with a user. For more information, see [Section 1.4.2.4, "Associating a WebLogic User with a Security Context."](#)

You can use the same properties on either a client or a server. If you define the properties on a server-side object, a local context is used. If you define the properties on a client or another WebLogic Server instance, the context delegates to a remote context running on the WebLogic Server instance specified by the `Context.PROVIDER_URL` property. A remote object bound to the server will not be serviced by `peerGone`, and will not be reachable if the client should fail.

There are some properties that cannot be changed after the creation of the context. These properties include provider URL, user credentials, and factories. `AddToEnvironment` can be used to change other properties after the creation of the context.

[Example 1–1](#) shows how to obtain a context using the properties `Context.INITIAL_CONTEXT_FACTORY` and `Context.PROVIDER_URL`.

Example 1–1 Obtaining a Context

```
Context ctx = null;
Hashtable ht = new Hashtable();
ht.put(Context.INITIAL_CONTEXT_FACTORY,
        "weblogic.jndi.WLInitialContextFactory");
ht.put(Context.PROVIDER_URL,
        "t3://localhost:7001");

try {
    ctx = new InitialContext(ht);
    // Use the context in your program
}
catch (NamingException e) {
    // a failure occurred
}
finally {
    try {ctx.close();}
    catch (Exception e) {
        // a failure occurred
    }
}
```

Additional WebLogic-specific properties are also available for controlling how objects are bound into the cluster-wide JNDI tree. Bindings may or may not be replicated across the JNDI tree of each server within the cluster due to the way these properties are set. Properties such as these are identified by constants in the `weblogic.jndi.WLContext` class. For more information about JNDI-related clustering issues, see [Section 1.4.6.5, "Using WebLogic JNDI from a Client in a Clustered Environment."](#)

1.4.2.1 Creating a Context Using a Hashtable

You can create a context with a hashtable in which you have specified the properties described in [Section 1.4.2, "Setting Up JNDI Environment Properties for the Initial Context."](#)

To do so, pass the hashtable to the constructor for `InitialContext`. The property `java.naming.factory.initial` is used to specify how the `InitialContext` is created. To use WebLogic JNDI, you must always set the `java.naming.factory.initial` property to `weblogic.jndi.WLInitialContextFactory`. This setting identifies the factory that actually creates the context.

1.4.2.2 Creating a Context Using a WebLogic Environment Object

You can also create a context by using a WebLogic environment object implemented by `weblogic.jndi.environment`. Although the environment object is WebLogic-specific, it offers the following advantages:

- A set of defaults which reduces the amount of code you need to write.
- Convenience `set()` methods that provide compile-time type-safety. The type-safety `set()` methods can save you time both writing and debugging code.

The WebLogic environment object provides the following defaults:

- If you do not specify an `InitialContext` factory, `WLInitialContextFactory` is used.
- If you do not specify a user and password in the `Context.SECURITY_PRINCIPAL` and `Context.CREDENTIALS` properties, the guest user and password are used unless the thread has already been associated with a user.
- If you do not specify a `Context.PROVIDER_URL` property, `t3://localhost:7001` is used.

If you want to create an `InitialContext` with these defaults, write the following code:

```
Environment env = new Environment();
Context ctx = env.getInitialContext();
```

If you want to set only a WebLogic Server to a Distributed Name Service (DNS) name for client cluster access, write the following code:

```
Environment env = new Environment();
env.setProviderURL("t3://myweblogiccluster.com:7001");
Context ctx = env.getInitialContext();
```

Note: Every time you create a new JNDI environment object, you are creating a new security scope. This security scope ends with a `context.close()` method. The `environment.getInitialContext()` method does not work correctly with the IIOP protocol.

[Example 1–2](#) illustrates using a JNDI environment object to create a security context.

Example 1–2 Creating a Security Context with a JNDI Environment Object

```
weblogic.jndi.Environment environment = new weblogic.jndi.Environment();
environment.setInitialContextFactory(
    weblogic.jndi.Environment.DEFAULT_INITIAL_CONTEXT_FACTORY);
environment.setProviderURL("t3://bross:4441");
environment.setSecurityPrincipal("guest");
```

```
environment.setSecurityCredentials("guest");
InitialContext ctx = environment.getInitialContext();
```

1.4.2.3 Creating a Context from a Server-Side Object

You may also need to create a context from an object (an Enterprise JavaBean (EJB) or Remote Method Invocation (RMI) object) that is instantiated in the Java Virtual Machine (JVM) of WebLogic Server. When using a server-side object, you do not need to specify the `Context.PROVIDER_URL` property. User names and passwords are required only if you want to sign in as a specific user.

To create a context from within a server-side object, you first must create a new `InitialContext`, as follows:

```
Context ctx = new InitialContext();
```

You do not need to specify a factory or a provider URL. By default, the context is created as a context and is connected to the local naming service.

1.4.2.4 Associating a WebLogic User with a Security Context

See [Section 1.4.2.5, "JNDI Contexts and Threads."](#)

1.4.2.5 JNDI Contexts and Threads

When you create a JNDI context with a user name and password, you associate a user with a thread. When the context is created, the user is pushed onto the context stack associated with the thread. Before starting a new context on the thread, you must close the first context so that the first user is no longer associated with the thread. Otherwise, users are pushed down in the stack each time a new context created. This is *not* an efficient use of resources and may result in the incorrect user being returned by `ctx.lookup()` calls. This scenario is illustrated by the following steps:

1. Create a second context (with user name and credential) called `ctx2` for `user2`. At this point, the thread has a stack of users associated with it. `User2` is at the top of the stack and `user1` is below it in the stack, so `user2` is used is the current user.
2. If you do a `ctx1.lookup("abc")` call, `user2` is used as the identity rather than `user1`, because `user2` is at the top of the stack. To get the expected result, which is to have `ctx1.lookup("abc")` call performed as `user1`, you need to do a `ctx2.close()` call. The `ctx2.close()` call removes `user2` from the stack associated with the thread and so that a `ctx1.lookup("abc")` call now uses `user1` as expected.

Note: When the `weblogic.jndi.enableDefaultUser` flag is enabled, there are two situations where a `close()` call does not remove the current user from the stack and this can cause JNDI context problems. For information on how to avoid JNDI context problems, see [Section 1.4.2.5.1, "How to Avoid Potential JNDI Context Problems."](#)

1.4.2.5.1 How to Avoid Potential JNDI Context Problems Issuing a `close()` call is usually as described in [Section 1.4.2.5, "JNDI Contexts and Threads."](#) However, the following is an exception to the expected behavior that occurs when the `weblogic.jndi.enableDefaultUser` flag is enabled:

1.4.2.5.2 Last Used When using IIOP, an exception to expected behavior arises when there is one context on the stack and that context is removed by a `close()`. The

identity of the last context removed from the stack determines the current identity of the user. This scenario is described in the following steps:

1. Do a `ctx1.close()` call.
2. Do a `ctx1.lookup()` call. *The current identity is user1.*
3. Create a context (with user name and credential) called `ctx2` for `user2`. In the process of creating the context, `user2` is associated with the thread and stored in the stack, that is, the current identity is set to `user2`.
4. Do a `ctx2.close()` call.
5. Do a `ctx2.lookup()` call. *The current identity is user2.*

1.4.3 Using the Context to Look Up a Named Object

The `lookup()` method on the context is used to obtain named objects. The argument passed to the `lookup()` method is a string that contains the name of the desired object. [Example 1–3](#) shows how to retrieve an EJB named `ServiceBean`.

Example 1–3 Looking Up a Named Object

```
try {
    ServiceBean bean = (ServiceBean)ctx.lookup("ejb.serviceBean");
} catch (NameNotFoundException e) {
    // binding does not exist
} catch (NamingException e) {
    // a failure occurred
}
```

1.4.4 Using a Named Object to Get an Object Reference

EJB client applications get object references to EJB remote objects from EJB Homes. RMI client applications get object references to other RMI objects from an initial named object. Both initial named remote objects are known to WebLogic Server as factories. A factory is any object that can return a reference to another object that is in the WebLogic namespace.

The client application invokes a method on a factory to obtain a reference to a remote object of a specific class. The client application then invokes methods on the remote object, passing any required arguments.

[Example 1–4](#) contains a code fragment that obtains a remote object and then invokes a method on it.

Example 1–4 Using a Named Object to Get an Object Reference

```
ServiceBean bean = ServiceBean.Home.create("ejb.ServiceBean")
Servicebean.additem(66);
```

1.4.5 Closing the Context

After clients finish working with a context, Oracle recommends that the client close the context in order to release resources and avoid memory leaks. Oracle recommends that you use a `finally{}` block and wrap the `close()` method in a `try{}` block. If you attempt to close a context that was never instantiated because of an error, the Java client application throws an exception.

In [Example 1–5](#), the client closes the context, releasing the resource being used.

Example 1–5 Closing the Context

```
try {
    ctx.close();
} catch () {
    //a failure occurred
}
```

1.4.6 Using WebLogic JNDI in a Clustered Environment

The intent of WebLogic JNDI is to provide a naming service for Java EE services, specifically EJB, RMI, and Java Messaging Service (JMS). Therefore, it is important to understand the implications of binding an object to the JNDI tree in a clustered environment.

The following sections discuss how WebLogic JNDI is implemented in a clustered environment and offer some approaches you can take to make your own objects available to JNDI clients.

1.4.6.1 Using the Relationship of RMI and JNDI to Enable WebLogic Clusters

WebLogic RMI is the enabling technology that allows clients in one JVM to access EJBs and JMS services from a client in another JVM. RMI stubs marshal incoming calls from the client to the RMI object. To make Java EE services available to a client, WebLogic binds an RMI stub for a particular service into its JNDI tree under a particular name. The RMI stub is updated with the location of other instances of the RMI object as the instances are deployed to other servers in the cluster. If a server within the cluster fails, the RMI stubs in the other server's JNDI tree are updated to reflect the server failure.

When a client connects to a cluster, it is actually connecting to one of the WebLogic Server instances in the cluster. Because the JNDI tree for this WebLogic Server instance contains the RMI stubs for all services offered by the other WebLogic Servers in the cluster in addition to its own services, the cluster appears to the client as one WebLogic Server instance hosting all of the cluster-wide services. When a new WebLogic Server instance joins a cluster, each WebLogic Server instance already in the cluster is responsible for sharing information about its own services to the new WebLogic Server instance. With the information collected from all the other servers in the cluster, the new server will create its own copy of the cluster-wide JNDI tree.

RMI stubs significantly affect how WebLogic JNDI is implemented in a clustered environment:

- RMI stubs are relatively small. This allows WebLogic JNDI to replicate stubs across all WebLogic Server instances in a cluster with little overhead in terms of server-to-server cross-talk.
- RMI stubs serve as the mechanism for replication across a cluster. An instance of a RMI object is deployed to a single WebLogic Server instance, however, the stub is replicated across the cluster.

1.4.6.2 Making Custom Objects Available to a WebLogic Server Cluster

When you bind a custom object (a non-RMI object) into a JNDI tree in a WebLogic Server cluster, the object is replicated across all the servers in the cluster. However, if the host server goes down, the custom object is removed from the cluster's JNDI tree. Custom objects are not replicated unless the custom object is bound again. You need to unbind and rebind a custom object every time you want to propagate changes made to the custom object. Therefore, WebLogic JNDI should not be used as a distributed object cache. You can use a third-party solution with WebLogic Server to provide distributed caches.

Suppose the custom object needs to be accessed only by EJBs that are deployed on only one WebLogic Server instance. Obviously it is unnecessary to replicate this custom object throughout all the WebLogic Server instances in the cluster. In fact, you should avoid replicating the custom object in order to avoid any performance degradation due to unnecessary server-to-server communication. To create a binding that is not replicated across WebLogic Server instances in a cluster, you must specify the `REPLICATE_BINDINGS` property when creating the context that binds the custom object to the namespace. [Example 1-6](#) illustrates the use of the `REPLICATE_BINDINGS` property.

Example 1-6 Using the `REPLICATE_BINDINGS` Property

```
Hashtable ht = new Hashtable();
//turn off binding replication
ht.put(WLContext.REPLICATE_BINDINGS, "false");
try {
    Context ctx = new InitialContext(ht);
    //bind the object
    ctx.bind("my_object", MyObject);
} catch (NamingException ne) {
    //failure occurred
}
```

When you are using this technique and you need to use the custom object, you must explicitly obtain an `InitialContext` for the WebLogic Server instance. If you connect to any other WebLogic Server instance in the cluster, the binding does not appear in the JNDI tree.

If you need a custom object accessible from any WebLogic Server instance in the cluster, deploy the custom object on each WebLogic Server instance in the cluster without replicating the JNDI bindings.

When using WebLogic JNDI to replicate bindings, the bound object will be handled as if it is owned by the host WebLogic Server instance. If the host WebLogic Server instance fails, the custom object is removed from all the JNDI trees of all WebLogic Server instances in the cluster. This behavior can have an adverse effect on the availability of the custom object.

1.4.6.3 Data Caching Design Pattern

A common task in Web applications is to cache data used by multiple objects for a period of time to avoid the overhead associated with data computation or connecting to another service.

Suppose you have designed a custom data caching object that performs well on a single WebLogic Server instance and you would like to use this same object within a WebLogic Server cluster. If you bind the data caching object in the JNDI tree of one of the WebLogic Server instances, WebLogic JNDI will, by default, copy the object to each of the other WebLogic Server instances in the cluster. It is important to note that since this is not an RMI object, what you are binding into the JNDI tree (and copying to the other WebLogic Server instances) is the object itself, not a stub that refers to a single instance of the object hosted on one of the WebLogic Server instances. Do not assume from the fact that WebLogic Server copies a custom object between servers that custom objects can be used as a distributed cache. Remember the custom object is removed from the cluster if the WebLogic Server instance to which it was bound fails and changes to the custom object are not propagated through the cluster unless the object is unbound and rebound to the JNDI tree.

For performance and availability considerations, it is often desirable to avoid using WebLogic JNDI's binding replication to copy large custom objects with high availability requirements to all of the WebLogic Server instances in a cluster. As an alternative, you can deploy a separate instance of the custom object on each of the WebLogic Server instances in the cluster. When binding the object to each WebLogic Server instance's JNDI tree, you should make sure to turn off binding replication as described in [Section 1.4.6.2, "Making Custom Objects Available to a WebLogic Server Cluster"](#). In this design pattern, each WebLogic Server instance has a copy of the custom object but you will avoid copying large amounts of data from server to server.

Regardless of which approach you use, each instance of the object should maintain its own logic for when it needs to refresh its cache independently of the other data cache objects in the cluster. For example, suppose a client accesses the data cache on one WebLogic Server instance. It is the first time the caching object has been accessed, so it computes or obtains the information and saves a copy of the information for future requests. Now suppose another client connects to the cluster to perform the same task as the first client only this time the connection is made to a different WebLogic Server instance in the cluster. If this is the first time this particular data caching object has been accessed, it will need to compute the information regardless of whether other data caching objects in the cluster already have the information cached. Of course, for any future requests, this instance of the data cache object will be able to refer to the information it has saved.

1.4.6.4 Exactly-Once-Per-Cluster Design Pattern

In some cases, it is desirable to have a service that appears only once in the cluster. This is accomplished by deploying the service on one machine only. For RMI objects, you can use the default behavior of WebLogic JNDI to replicate the binding (the RMI stub) and the single instance of your object will be accessible from all WebLogic Server instances in the cluster. This is referred to as a *pinned* service. For non-RMI objects, make sure that you use the `REPLICATE_BINDINGS` property when binding the object to the namespace. In this case, you will need to explicitly connect to the host WebLogic Server instance to access the object. Alternatively, you can create an RMI object that is deployed on the same host WebLogic Server instance that can act as a proxy for your non-RMI object. The stub for the proxy can be replicated (using the default WebLogic JNDI behavior) allowing clients connected to any WebLogic Server instance in the cluster to access the non-RMI object via the RMI proxy.

For services with high-availability requirements, you can configure automatic migration of an RMI object to another server. For more information about automatic migration, see "Whole Server Migration" in *Administering Clusters for Oracle WebLogic Server*.

1.4.6.5 Using WebLogic JNDI from a Client in a Clustered Environment

The JNDI binding for an object can appear in the JNDI tree for one WebLogic Server instance in the cluster, or it can be replicated to all the WebLogic Server instances in the cluster. If the object of interest is bound in only one WebLogic Server instance, you must explicitly connect to the host WebLogic Server instance by setting the `Context.PROVIDER_URL` property to the host WebLogic Server URL when creating the `InitialContext`, as described in [Section 1.4.1, "Using WebLogic JNDI to Connect a Java Client to a Single Server."](#)

In most cases, however, the object of interest is either a clustered service or a pinned service. As a result, a stub for the service is displayed in the JNDI tree for each WebLogic Server instance in the cluster. In this case, the client does not need to name a specific WebLogic Server instance to provide its naming service. In fact, it is best for the client to simply request that a WebLogic cluster provide a naming service, in which

case the context factory in WebLogic Server can choose whichever WebLogic Server instance in the cluster seems most appropriate for the client.

Currently, a naming service provider is chosen within WebLogic using a DNS name for the cluster that can be defined by the `ClusterAddress` attribute. This attribute defines the address to be used by clients to connect to a cluster. This address may be either a DNS host name that maps to multiple IP addresses or a comma separated list of single address host names or IP addresses. If network channels are configured, it is possible to set the cluster address on a per channel basis. See "Communications In a Cluster" in *Administering Clusters for Oracle WebLogic Server*.

The context that is returned to a client of clustered services is, in general, implemented as a failover stub that can transparently change the naming service provider if a failure (such as a communication failure) with the selected WebLogic Server instance occurs.

[Example 1-7](#) shows how a client uses the cluster's naming service.

Example 1-7 Using the Naming Service in a WebLogic Cluster

```
Hashtable ht = new Hashtable();
ht.put(Context.INITIAL_CONTEXT_FACTORY,
        "weblogic.jndi.WLInitialContextFactory");
ht.put(Context.PROVIDER_URL, "t3://acmeCluster:7001");
try {
    Context ctx = new InitialContext(ht);
    // Do the client's work
}
catch (NamingException ne) {
    // A failure occurred
}
finally {
    try {ctx.close();}
    catch (Exception e) {
        // a failure occurred
    }
}
```

The hostname specified as part of the provider URL is the DNS name for the cluster that can be defined by the `ClusterAddress` setting in a `Cluster` stanza of the `config.xml` file. `ClusterAddress` maps to the list of hosts providing naming service in this cluster. For more information, see "Understanding Cluster Configuration" in *Administering Clusters for Oracle WebLogic Server*.

In [Example 1-7](#), the cluster name `acmeCluster` is used to connect to any of the WebLogic Server instances in the cluster. The resulting context is replicated so that it can fail over transparently to any WebLogic Server instance in the cluster.

An alternative method of specifying the initial point of contact with the WebLogic Server cluster is to supply a comma-delimited list of DNS server names or IP addresses.

- The following example specifies a list of WebLogic Server instances using the same port:

```
ht.put(Context.PROVIDER_URL, "t3://acme1,acme2,acme3:7001");
```

All the WebLogic Server instances listen on the port specified at the end of the URL.

- The following example specifies a list of WebLogic Server instances using different ports:

```
ht.put(Context.PROVIDER_URL, "t3://node1:7001,node2:7002,node3:7003");
```

When you use a DNS name which maps to multiple servers, WebLogic Server relies on DNS for load balancing.

When you use a comma-delimited list of DNS names for WebLogic Server nodes, failover is accomplished using the round-robin method, with the request going to a randomly chosen server until that server fails to respond, after which the request will go to the next server on the list. This will continue for each server instance that fails.

Once the client has gotten a context, no additional load balancing occurs unless there is a failure, in which case a WebLogic Server instance will fail over to another node in the cluster.

A remote client will get the context from the first available server. A client that is local to a server in the cluster will never go to a remote server for JNDI operations.

When you look up a stub, the first invocation of the stub will ordinarily go to the server from which you got the context. If the stub is clusterable, subsequent invocations will be load balanced based on the user defined load balancing policy.

For additional information about JNDI and clusters see "Overview" in *Administering Clusters for Oracle WebLogic Server*.

1.4.7 Using JNDI from Within Java EE Components

Although it is possible for Java EE components to use the global environment directly, it is preferable to use the component environment. Each Java EE component within a Java EE application has its own component environment which is set up based on information contained in the component's deployment descriptors.

Java EE components are able to look up their component environments using the following code:

```
Context ctx = new InitialContext();
Context comp_env = (Context)ctx.lookup("java:comp/env");
```

Because you are working within a Java EE component, you do not need to set up the hashtable or environment objects to define the connection information.

This context is used in the same way as the global environment, however, the names you use are the ones defined in the deployment descriptor for your component. For example, if you have an `ejb-ref` in your deployment descriptor that looks like:

```
<ejb-ref>
...
<ejb-ref-name>ejb1</ejb-ref-name>
<ejb-ref-type>Session</ejb-ref-type>
<home>ejb1.EJB1Home</home>
<remote>ejb1.EJB1</remote>
...
</ejb-ref>
```

you would look up the name defined with the `<ejb-ref-name>` setting, which in this case is "ejb1".

Using the component environment rather than the global environment to set your JNDI name is advantageous because the name it refers to is resolved during deployment. This means that naming conflicts can be resolved without rewriting the code.

1.4.8 Setting Up Foreign JNDI

Foreign JNDI is an API that allows you to access objects on a remote JNDI tree without having to connect directly to the remote tree.

Foreign JNDI enables you to make links to a JNDI tree on another server or provider including, but not limited to, WebLogic Server, or a JNDI tree in a Java program. Once you have configured Foreign JNDI, you can use an object that is somewhere else with the same ease that you would use an object bound in your WebLogic Server instance.

To configure Foreign JNDI, create a `ForeignJNDIProvider` with the address of the remote JNDI provider whose objects you want to use, and create a user name and password to access those objects. Optionally, you can target Foreign JNDI references to specific servers, clusters, or both. (If no targets are selected, Foreign JNDI references will be deployed to the entire domain). Then you can create `ForeignJNDILinks` and `ForeignJNDIObjects` that set up a relationship between a name in the local JNDI tree to the object in the remote tree.

For more information on how to configure Foreign JNDI, see "Create a foreign JNDI provider" in the *Oracle WebLogic Server Administration Console Online Help*.