# Oracle® Internet Directory

Administrator's Guide

Release 2.1.1

September 2000

Part No. A86101-01

ORACLE®

Oracle Internet Directory Administrator's Guide, Release 2.1.1

Part No. A86101-01



Primary Author: Richard Smith

Contributing Authors: Deborah Steiner, Sandy Venning

Contributors: Tridip Bhattachrya, Margaret Chou, Raj Gupta, Ashish Kolli, Stephen Lee, Michael Mesaros, Radikah Moolky, Olaf Stullich, David Saslav, Hari Sastry, Gurudat Shakshikumar, Amit Sharma, Daniel Shih, Saurabh Shrivastava, Uppili Srinivasan



# Contents

## 2 Concepts and Architecture

## 3 Preliminary Tasks

## 6 Managing the Directory Schema

## 7 Managing Directory Entries

## 8 Managing Secure Sockets Layer (SSL)



## 9 Managing Directory Access Control

## 10 Managing Directory Replication

## 11 Synchronizing with Multiple Directories



## 12 Managing National Language Support (NLS)

## 14 Capacity Planning



## 15 Tuning

## 16 High Availability And Failover



## Part IV Appendixes

## A Syntax for LDIF and Command Line Tools

## D Using Access Control Directive Format



## E Schema Elements

# Send Us Your Comments

**Oracle Internet Directory Administrator's Guide, Release 2.1.1**

**Part No. A86101-01**

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the chapter, section, and page number (if available). You can send comments and suggestions about this manual to the Information Development department at the following e-mail address:

- E-mail - infodev@us.oracle.com
- FAX - (650) 506-7228. Attn: Oracle Internet Directory Documentation Manager
- Postal service:
  Oracle Corporation
  Oracle Internet Directory Documentation Manager
  500 Oracle Parkway, 4op7
  Redwood Shores, CA 94065
  U.S.A.

If you would like a reply, please give your name, address, and telephone number.

If you have problems with the software, please contact your local Oracle Support Services.

# Preface

*Oracle Internet Directory Administrator's Guide* describes the features, architecture, and administration of Oracle Internet Directory. For information about installation, see the installation documentation for your operating system.

## Audience

This book is intended for anyone who performs system administration tasks for the Oracle Internet Directory. You should be familiar with either the UNIX operating system or the Microsoft Windows NT operating system in order to understand the line-mode commands and examples. You can perform all of the tasks through the line-mode commands, and you can perform most of the tasks through Oracle Directory Manager, which is operating system-independent.

## Organization

This book contains the chapters and appendixes listed in this section. Oracle Corporation encourages you to read the conceptual and other introductory material presented in Part I before attempting installation and maintenance.

### Part I: Getting Started

Part I provides an overview of the product and its features, a conceptual foundation necessary to configure and manage a directory, instructions for starting a directory server, and an introduction to the various administration tools. Specific chapters and their descriptions are:

| | |
|---|---|
| Chapter 1, "Introduction" | Provides an introduction to directories, LDAP, and Oracle Internet Directory features. |
| Chapter 2, "Concepts and Architecture" | Gives an overview of online directories and Lightweight Directory Access Protocol (LDAP). Provides conceptual descriptions of directory entries, attributes, object classes, naming contexts, schemas, distributed directories, security, and National Language Support. This chapter also discusses Oracle Internet Directory architecture. |
| Chapter 3, "Preliminary Tasks" | Discusses how to prepare your directory for configuration and use. It tells you how to start and stop OID Monitor and instances of Oracle directory server and Oracle directory replication server. It discusses the need to reset the default security configuration. Finally, it discusses how to upgrade from earlier releases of Oracle Internet Directory, and how to migrate data from other LDAP-compliant directories. |
| Chapter 4, "Using the Administration Tools" | Explains how to use the various administration tools: Oracle Directory Manager, command line tools, bulk tools, Catalog Management tool, OID Database Password Utility, replication tools, and Database Statistics Collection tool |

### Part II: Managing Oracle Internet Directory

Part II guides you through the tasks required to configure and maintain Oracle Internet Directory. Specific chapters and their descriptions are:

| | |
|---|---|
| Chapter 5, "Managing an Oracle Directory Server" | Provides instructions for managing server configuration set entries; setting system operational attributes; managing naming contexts and password encryption; configuring searches; managing super, guest, and proxy users; setting debug logging levels; using audit log; viewing active server instance information; and changing the password to an Oracle database server. |
| Chapter 6, "Managing the Directory Schema" | Explains what a directory schema is, what an object class is, and what an attribute is. It tells you how to manage the Oracle Internet Directory schema by using Oracle Directory Manager and the command line tools. |
| Chapter 7, "Managing Directory Entries" | Explains how to search, view, add, modify and manage entries by using Oracle Directory Manager and the command line tools. |
| Chapter 8, "Managing Secure Sockets Layer (SSL)" | Introduces and explains how to configure the features of Secure Sockets Layer (SSL). |
| Chapter 9, "Managing Directory Access Control" | Provides an overview of access control policies and describes how to administer directory access. |
| Chapter 10, "Managing Directory Replication" | Explains replication; how to install and initialize Oracle directory replication server software the first time, and how to install new nodes into an environment where that software is already installed. |
| Chapter 11, "Synchronizing with Multiple Directories" | Explains how synchronization takes place between Oracle Internet Directory and other directories. It tells you how to enable other directories to synchronize with Oracle Internet Directory. |
| Chapter 12, "Managing National Language Support (NLS)" | Discusses National Language Support (NLS) as used by Oracle Internet Directory. |

### Part III: Deploying Oracle Internet Directory

Part III discusses deployment considerations. Specific chapters and their descriptions are:

| | |
|---|---|
| Chapter 13, "Deployment Considerations" | Discusses issues to consider when deploying Oracle Internet Directory. This chapter helps you assess the requirements of a directory in an enterprise and make effective deployment choices. |
| Chapter 14, "Capacity Planning" | Gives guidelines for capacity planning for an Oracle Internet Directory installation. |
| Chapter 15, "Tuning" | Gives guidelines for tuning an Oracle Internet Directory installation. |
| Chapter 16, "High Availability And Failover" | Discusses the high availability and failover features and deployment guidelines for Oracle Internet Directory. |

### Part IV: Appendixes

| | |
|---|---|
| Appendix A, "Syntax for LDIF and Command Line Tools" | Provides syntax, usage notes, and examples for LDAP Data Interchange Format and LDAP command line tools. |
| Appendix B, "Adding a DSA Using the Database Copy Procedure" | Describes an alternate method of adding a node to a replicated directory system if the directory is very large. |
| Appendix C, "Using Oracle Wallet Manager" | Describes and explains how to use Oracle Wallet Manager to create and manage wallets and certificates. |
| Appendix D, "Using Access Control Directive Format" | Describes the format (syntax) of Access Control Information Items(ACIs). |
| Appendix E, "Schema Elements" | Lists schema elements supported in Oracle Internet Directory. |
| Appendix F, "Migrating Data from Other LDAP-Compliant Directories" | Explains the steps to migrate data from LDAP v3-compatible directories into Oracle Internet Directory. |
| Appendix G, "Troubleshooting" | Lists possible failures and error codes and their probable causes |

## Related Documentation

For related Oracle information, refer to the following:

- Online help available through Oracle Directory Manager
- Oracle8*i* documentation set

For more information about concepts discussed in this book, see the books and online articles below. Most of these entries also contain references to other publications.

Chadwick, David. *Understanding X.500 The Directory*. Thomson Computer Press, 1996. This book is now out of print, but is available online at: http://www.salford.ac.uk/its024/Version.Web/Contents.htm

Hodges, Jeff, Staff Scientist, Oblix, Inc., http://www.kingsmountain.com/ldapRoadmap.shtml

Howes, Tim and Mark Smith. *LDAP: Programming Directory-enabled Applications with Lightweight Directory Access Protocol*. Macmillan Technical Publishing, 1997.

Howes, Tim, Mark Smith and Gordon Good, *Understanding and Deploying LDAP Directory Services*. Macmillan Technical Publishing, 1999.

Kosiur, Dave, LDAP: "The next-generation directory?," *SunWorld Online*, October 1997.

Radicati, Sara, *X.500 Directory Services, Technology and Deployment*, International Thomson Computer Press, 1994.

*University of Michigan LDAP Repository*, http://www.umich.edu/~dirsvcs/ldap/index.html

## Conventions

This manual uses the following conventions:

| Convention | Meaning |
|---|---|
| .<br>.<br>. | Vertical ellipsis points in an example mean that information not directly related to the example has been omitted. |
| . . . | Horizontal ellipsis points in statements or commands mean that parts of the statement or command have been omitted. |

| Convention | Meaning |
|---|---|
| **bold** | Boldface text indicates a term defined in the glossary, text you must type in a command, or a subheading. |
| *italics* | Italics indicate:<br>■ In a code example, a variable for which you must supply a value<br>■ In regular text, special emphasis<br>■ Book titles |
| `courier` | Courier is used for user input and code examples. |
| *syntax* | This typeface is used for syntax explanations in code examples. |
| < > | In code examples, angle brackets may enclose user-supplied names. |
| [ ] | Brackets enclose a choice of optional items from which you can choose one or none. |
| { } | Braces enclose a choice of required items from which you can choose one. |

# What's New in Oracle Internet Directory?

Oracle Internet Directory release 2.1.1 contains these new features:

| Feature | Information |
|---|---|
| Attribute options, including language codes | "Attribute Options" on page 2-7 for a conceptual discussion<br>"Managing Entries with Attribute Options" on page 7-17 |
| Change log purging enhancements | "Change Log Purging" on page 2-30 for a conceptual discussion<br>"Oracle Directory Replication Server Parameters" on page 10-10 |
| Enhanced support for these operational attributes:<br>▪ `creatorsName`<br>▪ `createTimestamp`<br>▪ `modifiersName`<br>▪ `modifyTimestamp` | "Kinds of Attribute Information" on page 2-5 for a conceptual discussion<br>"Setting System Operational Attributes" on page 5-13<br>"Example 6: Searching for All User Attributes and Specified Operational Attributes" on page A-20 for an example of a search operation using the `createTimestamp` attribute |
| Human intervention queue manipulation tool | "Using the Replication Tools" on page 4-15 for a brief explanation of this tool<br>"Using the Human Intervention Queue Manipulation Tool" on page 10-30 |
| Migration from other LDAP-compliant directories | Appendix F, "Migrating Data from Other LDAP-Compliant Directories" |
| Object class explosion | "Guidelines for Adding Object Classes" on page 6-3 for an explanation of how to use this feature when adding object classes |

| Feature | Information |
|---|---|
| OID Database Statistics Collection Tool | "Using the OID Database Statistics Collection Tool" on page 4-15 |
| Password encryption enhancements | "Password Encryption" on page 2-17 for a conceptual discussion<br>"Managing Password Encryption" on page 5-17 for instructions on setting password encryption |
| OID reconciliation tool | "Using the Replication Tools" on page 4-15 for a brief explanation of this tool<br>"Using the OID Reconciliation Tool" on page 10-33 |
| Replication node deletion | "Deleting a Replication Node" on page 10-25 |
| Synchronization with multiple directories in a metadirectory environment | "Synchronizing with Other Directories in a Metadirectory Environment" on page 2-46 for a conceptual discussion<br>Chapter 11, "Synchronizing with Multiple Directories" |
| Upgrade procedures | "Upgrading from an Earlier Release of Oracle Internet Directory" on page 3-9 |

# Part I

## Getting Started

Part I provides an overview of Oracle Internet Directory and its features, a conceptual foundation necessary to correctly configure and manage a directory, specific instructions on how to get started, and an introduction to the administration tools.

This part contains these chapters:

# 1

# Introduction

This chapter is an introduction to online directories, provides an overview of the Lightweight Directory Application Protocol (LDAP) version 3, and explains some of the unique features and benefits of Oracle Internet Directory.

This chapter contains these topics:

# What Is a Directory?

A directory organizes information so that you can find it easily. It lists objects—for example, people, books in a library, or merchandise in a department store—and gives details about each one. A telephone book is a familiar type of directory, a card catalog in a library is another, and a department store catalog still another.

This section contains these topics:



## Online Directories

An online directory is a specialized database that stores and retrieves collections of information about objects. Such information can represent any resources that require management: employee names, titles, and security credentials; information about e-commerce partners; or information about shared network resources such as conference rooms and printers.

Directories can be used by a variety of users and applications, for a variety of purposes. A few typical scenarios include:

- An employee searching for corporate whitepage information, and, through a mail client, looking up email addresses
- An application, such as a message transport agent, locating a user's mail server
- A database application identifying user role information

## The Difference Between Online Directories and Relational Databases

A database is a structured collection of data. Although an online directory is a database, it is not a **relational database**. The following table contrasts online directories with relational databases.

| Online Directories | Relational Databases |
|---|---|
| **Primarily read-focused.** Typical use involves a relatively small number of data updates, and a potentially large number of data retrievals. | **Primarily write-focused.** Typical use involves continuous recording of transactions, with retrievals done relatively infrequently. |
| **Designed to handle relatively simple transactions on relatively small units of data.** For example, an application might use a directory simply to store and retrieve an e-mail address, a telephone number, or a digital portrait. | **Designed to handle large and diverse transactions using many operations on large units of data.** |
| **Designed to be location-independent.** Directory applications expect, at all times, to see the same information throughout the deployment environment—regardless of which server they are querying. If a queried server does not store the information locally, then it must either retrieve the information or point the client application to it transparently. | **Typically designed to be location-specific.** While a relational database can be distributed, it usually resides on a particular database server. |
| **Designed to store information in entries.** These entries might represent any resource customers wish to manage: employees, e-commerce partners, conference rooms, or shared network resources such as printers. Associated with each entry is a number of attributes, each of which may have one or more values assigned. For example, typical attributes for a `person` entry might include first and last names, e-mail addresses, the address of a preferred mail server, passwords or other login credentials, or a digitized portrait. | **Designed to store information as records in relational tables.** |

## The Problem: Multiple Special Purpose Directories

According to some estimates, each of the world's largest companies has an average of 180 different directories, each designated for a special purpose. Add to that the various enterprise applications, each with its own additional directory of user names, and the actual number of special purpose directories becomes even higher.

Managing so many special purpose directories can cause three problems:

- Inconsistent data: Updated information in one directory is not shared with all the other directories.
- Redundancy: The same information is represented in many different places in the enterprise.
- High cost of administration: Administrators must maintain essentially the same information in many different places.

For example, when an enterprise hires a new employee, administrators must create a new user identity on the network, create a new e-mail account, add the user to the human-resources database, and set up all applications that the employee may need—for example, user accounts on development, testing, and production database systems. Later, if the employee leaves the company, administrators must reverse the process to disable all these user accounts. In addition to this administrative overhead, it can be difficult for multiple administrators entering redundant information in multiple systems to synchronize this employee information across all systems. The result can be inconsistent data across the enterprise.

## The Solution: The LDAP-Compliant General Purpose Directory

Clearly there is need for a more general purpose directory infrastructure, one based on a common standard for supporting a wide variety of applications and services. Oracle Internet Directory answers this need through its use of the **Lightweight Directory Access Protocol (LDAP)**.

# What Is LDAP?

The **Lightweight Directory Access Protocol (LDAP)** is a standard, extensible directory access protocol. It is a common language that LDAP clients and servers use to communicate.

LDAP was conceived as an internet-ready, lightweight implementation of the International Standardization Organization (ISO) X.500 standard for directory services. It requires a minimal amount of networking software on the client side, which makes it particularly attractive for internet-based, thin client applications.

This section contains these topics:



## LDAP and Simplified Directory Management

The LDAP standard simplifies management of directory information in three ways:

- It provides all users and applications in the enterprise with a single, well-defined, standard interface to a single, extensible directory service. This makes it easier to rapidly develop and deploy directory-enabled applications.
- It reduces the need to enter and coordinate redundant information in multiple services scattered across the enterprise.
- Its well-defined protocol and array of programmatic interfaces make it more practical to deploy internet-ready applications that leverage the directory.

## LDAP Version 3

The most recent version of LDAP, Version 3, was approved as a proposed Internet Standard by the Internet Engineering Task Force in December 1997. LDAP Version 3 improves on LDAP Version 2 in several important areas:

- National Language Support: LDAP Version 3 allows servers and clients to support characters used in every language in the world.
- Knowledge references (also called referrals): LDAP Version 3 implements a referral mechanism that allows servers to return references to other servers as a result of a directory query. This makes it possible to partition a **directory information tree (DIT)** (described in Chapter 2) across multiple LDAP servers, enabling global deployment.

- Security: A standard mechanism for supporting Simple Authentication and Security Layer (SASL) and Transport Layer Security (TLS) were added, providing LDAP with a comprehensive and extensible framework for data security.
- Extensibility: LDAP Version 3 enables vendors to extend existing LDAP operations through the use of mechanisms called controls.
- Feature and schema discovery: LDAP Version 3 enables publishing information useful to other LDAP servers and clients, such as the supported LDAP protocols and a description of the directory schema.

**See Also:**

- RFCs (Requests for Comments) 2251-2256 of the IETF, available on the Worldwide Web at:

  http://www.ietf.org/rfc.html

- "Related Documentation" on page xxvii for an additional list of resources on LDAP

# What Is Oracle Internet Directory?

Oracle Internet Directory is a general purpose directory service that enables retrieval of information about dispersed users and network resources. It combines **LDAP** Version 3 with the high performance, scalability, robustness, and availability of Oracle8*i*.

This section contains these topics:



## Oracle Internet Directory and Oracle8*i*

Oracle Internet Directory runs as an application on Oracle 8*i*. It communicates with the database, which may be on the same or on a different operating system, by using **Net8**, Oracle's operating system-independent database connectivity solution. Figure 1–1 illustrates this relationship.

***Figure 1–1 Oracle Internet Directory Architecture***

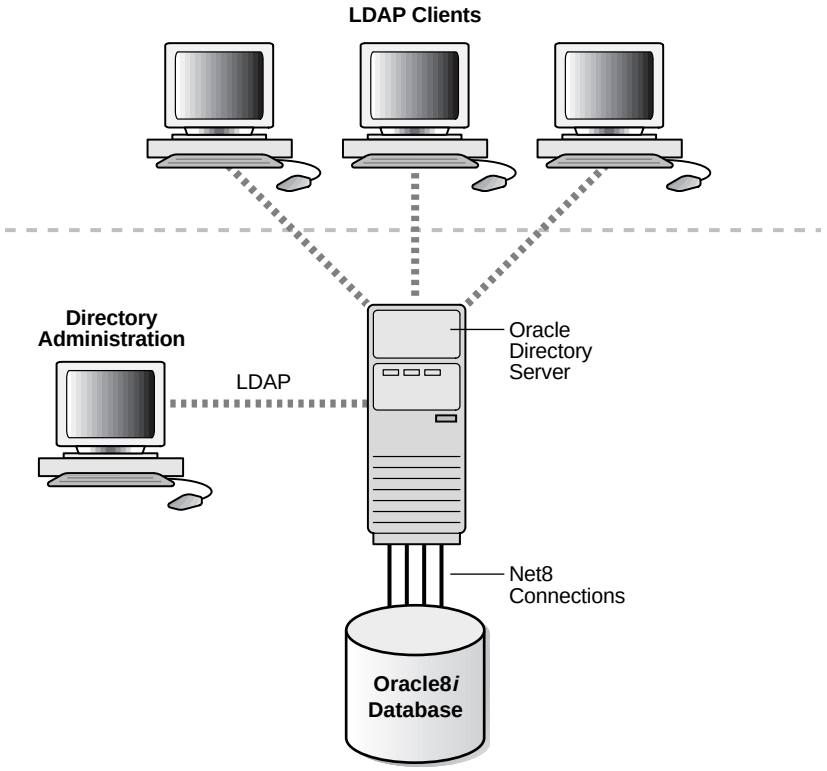


## Oracle Internet Directory Components

Oracle Internet Directory includes:

- Oracle directory server, which responds to client requests for information about people and resources, and to updates of that information, using a multi-tiered architecture directly over TCP/IP
- Oracle directory replication server, which replicates LDAP data between Oracle directory servers
- Oracle Directory Manager, a Java-based graphical user interface administration tool
- A variety of command line administration and data management tools

## The Advantages of Oracle Internet Directory

Oracle Internet Directory provides these significant benefits:

- Scalability
- High Availability
- Security

### Scalability

Oracle Internet Directory exploits the strengths of Oracle8*i*, enabling support for terabytes of directory information. In addition, such technologies as multithreaded LDAP servers and database connection pooling allow it to support thousands of concurrent clients with subsecond search response times.

Oracle Internet Directory also provides data management tools, such as Oracle Directory Manager and a variety of command line tools, for manipulating large volumes of LDAP data.

### High Availability

Oracle Internet Directory is designed to meet the needs of a variety of important applications. For example, Oracle Internet Directory supports full, multi-master replication between directory servers: If one server in a replication community becomes unavailable, then a user can access the data from another server. Information about changes made to data on a server is stored in special tables on the Oracle8*i* database. These are replicated throughout the directory environment by Oracle's **Advanced Symmetric Replication (ASR)**, a robust replication mechanism.

Oracle Internet Directory also takes advantage of all the availability features of the Oracle8*i*. Because directory information is stored securely in the Oracle8*i* database, it is protected by Oracle's backup capabilities. Additionally, the Oracle8*i* database, running with large datastores and heavy loads, can recover from system failures quickly.

### Security

Oracle Internet Directory offers comprehensive and flexible access control. An administrator can grant or control access to a specific directory object or to an entire directory subtree. Moreover, Oracle Internet Directory implements three levels of user authentication, namely, anonymous, password-based, and certificate-based using **Secure Socket Layer (SSL)** Version 3 for authenticated access and data privacy.

# 2

# Concepts and Architecture

This chapter provides conceptual descriptions of the basic elements of Oracle Internet Directory and discusses Oracle Internet Directory architecture.

This chapter contains these topics:

- Entries
- Attributes
- Object Classes
- Naming Contexts
- The Directory Schema
- Security
- National Language Support
- Oracle Internet Directory Architecture
- Distributed Directories: An Overview
- Distributed Directories: Replication
- Distributed Directories: Partitioning
- Synchronizing with Other Directories in a Metadirectory Environment

> **See Also:** "Related Documentation" on page xxvii for suggestions on further reading about LDAP-compliant directories

# Entries

In a directory, each collection of information about an object is called an **entry**. For example, a typical telephone directory includes entries for people, and a library card catalog contains entries for books. Similarly, an online directory might include entries for employees, conference rooms, e-commerce partners, or shared network resources such as printers.

Each entry in a directory is uniquely identified by a **distinguished name (DN)**. The distinguished name tells you exactly where the entry resides in the directory's hierarchy. This hierarchy is represented by a **directory information tree (DIT)**.

To understand the relation between a distinguished name and a directory information tree, look at the example in Figure 2–1.

***Figure 2–1 A Directory Information Tree***

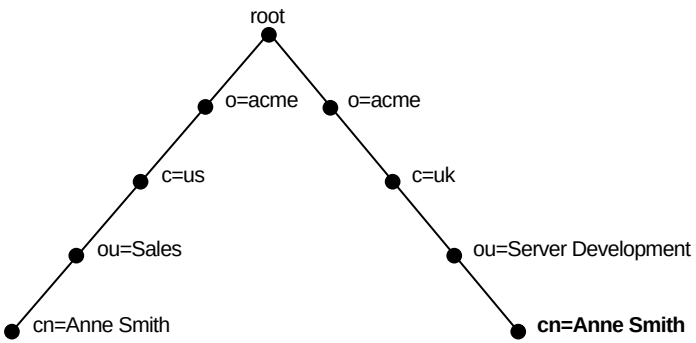


The DIT in Figure 2–1 diagrammatically represents entries for two employees of Acme Corporation who are both named Anne Smith. It is structured along geographical and organizational lines. The Anne Smith represented by the left branch works in the Sales division in the United States, while the other works in the Server Development division in the United Kingdom.

The Anne Smith represented by the right branch has the common name (cn) Anne Smith. She works in an organizational unit (ou) named Server Development, in the country (c) of Great Britain (uk), in the organization (o) Acme.

The DN for this "Anne Smith" entry is:

```
cn=Anne Smith,ou=Server Development,c=uk,o=acme
```

Note that the conventional format of a distinguished name places the lowest DIT component at the left, then follows it with the next highest component, thus moving progressively up to the root.

Within a distinguished name, the lowest component is called the **relative distinguished name (RDN)**. For example, in the above entry for Anne Smith, the RDN is `cn=Anne Smith`. Similarly, the RDN for the entry immediately above Anne Smith's RDN is `ou=Server Development`, the RDN for the entry immediately above `ou=Server Development` is `c=uk`, and so on. A DN is thus a sequence of RDNs separated by commas.

To locate a particular entry within the overall DIT, a client uniquely identifies that entry by using the full DN—not simply the RDN—of that entry. For example, within the global organization in Figure 2–1, to avoid confusion between the two Anne Smiths, you would use each one's full DN. (If there are potentially two employees with the same name in the same organizational unit, you could use additional mechanisms, such as identifying each employee with a unique identification number.)

**See Also:** Chapter 7, "Managing Directory Entries."

## Attributes

In a typical telephone directory, an **entry** for a person contains such information items as an address and a phone number. In an online directory, these information items are called **attributes**. Attributes in a typical employee entry can include, for example, a job title, an e-mail address, or a phone number.

For example, in Figure 2–2, the entry for Anne Smith in Great Britain (uk) has several attributes, each providing specific information about her. These are listed in the balloon to the right of the tree, and they include `emailaddrs`, `printername`, `jpegPhoto`, and `app preferences`. Moreover, each bullet in Figure 2–2 is also an entry with attributes, although the attributes for each are not shown.

***Figure 2–2 Attributes of the Entry for Anne Smith***

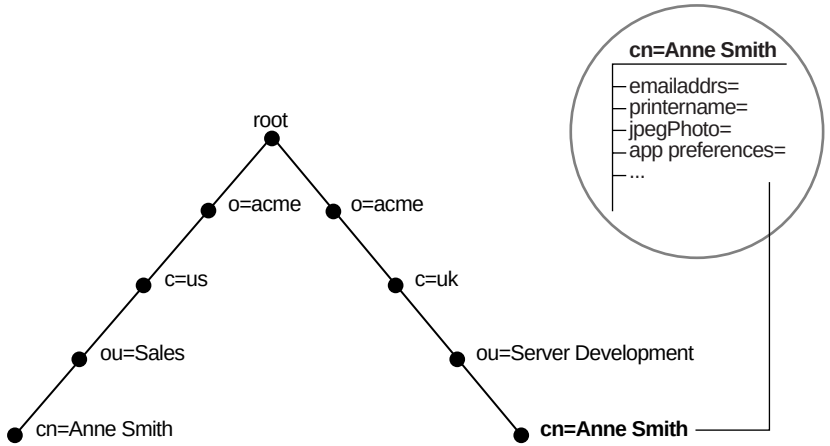


Each attribute consists of an attribute type and one or more attribute values. The attribute type identifies the kind of information that the attribute contains—for example, `jobTitle`. The attribute value is the particular occurrence of information appearing in that entry. For example, the value for the `jobTitle` attribute could be `manager`.

This section contains these topics:

- Kinds of Attribute Information
- Single-Valued and Multi-Valued Attributes
- Attribute Options
- Common LDAP Attributes
- Attribute Syntax
- Attribute Matching Rules

## Kinds of Attribute Information

Attributes contain two kinds of information.

| | |
|---|---|
| Application Information | This information is maintained and retrieved by the directory clients and is unimportant to the operation of the directory. A telephone number, for example, is application information. |
| Operational Information | This information pertains to the operation of the directory itself. Some operational information is specified by the directory to control the server—for example, the time stamp for the creation or modification of an entry, or the name of the user who creates or modifies an entry. Other operational information, such as access information, is defined by administrators and is used by the directory program in its processing. |

Any given attribute can hold either application information, or operational information, but not both.

To enhance your ability to search for entries, Oracle Internet Directory automatically creates several system operational attributes when you add an entry to the directory. These include:

| | |
|---|---|
| `creatorsName` | Name of the person creating the entry |
| `createTimestamp` | Time of entry creation in **UTC (Coordinated Universal Time)** |
| `modifiersName` | Name of person creating the entry |
| `modifyTimestamp` | Time of entry creation in UTC |

Moreover, when a user modifies an entry, Oracle Internet Directory automatically updates the `modifiersName` and `modifyTimestamp` attributes as follows:

| | |
|---|---|
| `modifiersName` | Name of person modifying the entry |
| `modifyTimestamp` | Time of entry modification in UTC |

**See Also:** "Setting System Operational Attributes" on page 5-13 for instructions on configuring system operational attributes

## Single-Valued and Multi-Valued Attributes

Attributes can be either single-valued or multi-valued. Single-valued attributes carry only one value in the attribute, whereas multi-valued attributes can have several. An example of a multi-valued attribute is a group membership list with names of everyone in the group.

## Common LDAP Attributes

Oracle Internet Directory implements all of the standard LDAP attributes. Table 2–1 shows some of the more common LDAP attributes.

***Table 2–1 Common LDAP Attributes***

| Attribute Type | Attribute String | Description | Example of Attribute Value |
|---|---|---|---|
| commonName | `cn` | Common name of an entry | `cn=Anne Smith` |
| domainComponent | `dc` | Component in a Domain Name System (DNS) | The following DN: `dc=uk,dc=acme,dc=com` |
| jpegPhoto | `jpegPhoto` | Photographic image in JPEG format. The path and file name of the JPEG image you want to include as an entry attribute. | `/photo/audrey.jpg` |
| organization | `o` | Name of an organization | `o=acme` |
| organizationalUnitName | `ou` | Name of a unit within an organization | `ou=Server Development` |
| owner | `owner` | Distinguished name of the person who owns the entry | The following line in an LDIF file: `owner: cn=Anne Smith, ou=Server Development, o= Acme, c=uk` |
| surname, sn | `sn` | Last name of a person | `Smith` |
| telephoneNumber | `telephoneNumber` | Telephone number | `telephoneNumber=(650) 123-4567` or `telephoneNumber=6501234567` |

**See Also:** Appendix E for a list of several proprietary attributes Oracle Internet Directory provides.

## Attribute Syntax

Attribute syntax is the format of the data that can be loaded into each attribute. For example, the syntax of the `telephoneNumber` attribute might require a telephone number to be a string of numbers containing spaces and hyphens. However, the syntax for another attribute might require specifying whether the data has to be in the form of a date, or whether the data can consist of numbers only. Each attribute must have one and only one syntax attached to it.

Oracle Internet Directory recognizes most of the syntax specified in RFC 2252, allowing you to associate most of the syntax described in that document with an attribute. In addition, Oracle Internet Directory also enforces some LDAP syntax. You cannot add new syntaxes beyond those already supported by Oracle Internet Directory.

**See Also:** "LDAP Syntax" on page E-7

## Attribute Matching Rules

In response to most incoming client requests, the directory server performs search and compare operations. During these operations, the directory server consults relevant **matching rule** to determine equality between the attribute value sought and the attribute value stored. For example, matching rules associated with the `telephoneNumber` attribute could cause "(650) 123-4567" to be matched with either "(650) 123-4567" or "6501234567" or both. When you create an attribute, you associate a matching rule with it.

Oracle Internet Directory implements all the standard LDAP matching rules. You cannot add new matching rules beyond those already supported by Oracle Internet Directory.

**See Also:** "Matching Rules" on page E-10

## Attribute Options

An attribute type can have various options that enable you to specify how the value for that attribute is made available in a search or a compare operation. For example, suppose that an employee has two addresses, one in London, the other in New York. Options for that employee's `address` attribute could allow you to store both addresses. Moreover, attribute options can include language codes. For example,

options for John Doe's `givenName` attribute could enable you to store his given name in both French and Japanese.

An attribute with one or more options has the properties—for example, matching rules and syntax— of its base attribute, which has no options. For example, suppose that `cn` is the base attribute. If `cn;lang-fr=Jean` is the French value for that base attribute, then it has the same matching rules and syntax as `cn`.

---

**Note:** You cannot use an attribute option within a DN. For example, the following DN is incorrect: `cn;lang-fr=Jean, ou=sales,o=acme,c=uk`.

---

**See Also:** Chapter 6, "Managing the Directory Schema."

## Object Classes

An **object class** is a group of **attributes**. When you define a directory **entry**, you assign one or more object classes to it. These object classes contain attributes, some of which are mandatory and some of which are optional.

For example, the `organizationalPerson` object class includes the mandatory attributes `commonName` (`cn`) and `surname` (`sn`). When you define an entry by using the `organizationalPerson` object class, you must specify values for these attributes. This object class also includes several optional attributes, including `telephoneNumber`, `uid`, `streetAddress`, and `userPassword`. You do not need to provide values for these latter attributes when using the `organizationalPerson` object class.

At installation, Oracle Internet Directory provides standard LDAP object classes, as well as several proprietary object classes. You cannot add mandatory attributes to the sets of attributes belonging to these predefined object classes. If a given object class does not contain all the attributes that you want for an entry, then you can do one of the following:

- Add optional attributes to an existing object class
- Define a new (base) object class
- Define an object subclass

**See Also:** Appendix E for a list of object classes in the schema installed with Oracle Internet Directory

This section contains these topics:



## Subclasses, Superclasses, and Inheritance

A **subclass** is an object class derived from another object class. The object class from which it is derived is called its **superclass**. For example, the object class `organizationalPerson` is a subclass of the object class `person`. Conversely, the object class `person` is the superclass of the object class `organizationalPerson`.

A subclass **inherits** all of the attributes belonging to its superclass. Entries may inherit the attributes defined by multiple object classes.

---

**Note:** In itself, an object class contains no values. Only an instance of an object class contains values. When a subclass inherits attributes from a superclass, it inherits only the attribute framework—not the attribute values—of the superclass

---

One special object class, called `top`, has no superclasses. It is one of the superclasses of every structural object class in the directory, and its attributes are inherited by every entry.

## Object Class Types

There are three types of object classes:

- Abstract
- Structural
- Auxiliary

### Abstract Object Classes

An abstract object class is a virtual object class, which means that it cannot be the only object class for an entry. For example, the object class `top` is an abstract object class. It is required as a superclass for all structural object classes, but it cannot be used alone.

The `top` object class includes the mandatory attribute `objectClass` as well as several optional attributes. The following list contains the names of the optional

attributes in `top` and either describes or points to further information about each one.

- orclGuid—Global identification which remains constant if the entry is moved
- creatorsName—See the appropriate IETF documentation.
- createTimestamp—See the appropriate IETF documentation.
- orclACI—See "orclACI" on page 9-2
- orclEntryLevelACI—See "orclEntryLevelACI" on page 9-3.

### Structural Object Classes

Structural object classes describe the basic aspects of an object. Most of the object classes that you use are structural object classes, and every entry should belong to at least one structural object class. Examples of structural object classes are `person` and `groupOfNames`.

These object classes place restrictions on which kinds of object classes can be created under any given object class. For example, a structure rule might require all objects below the `organization` (o) object class to be `organizational units` (ou). Following this rule, you could not enter `person` objects directly below an `organization` object class.

### Auxiliary Object Classes

Auxiliary object classes are groupings of attributes that expand the existing list of attributes in an entry. For example, suppose you have defined an entry as a member of two object classes, and you want to assign to that entry additional attributes that do not belong to either of the two object classes. You can create a new auxiliary object class containing the extra attributes, and then associate that auxiliary object class with the entry. This is an alternative to redefining existing object classes.

Unlike structural object classes, auxiliary classes do not place restrictions on where an entry may be stored.

---

**Note:** Oracle Internet Directory does not enforce structure rules. It therefore handles both structural and auxiliary object classes in the same way.

---

**See Also:** Chapter 6, "Managing the Directory Schema."

# Naming Contexts

A **naming context** is a subtree that resides entirely on one server. It must be contiguous, that is, it must begin at an **entry** that serves as the top of the subtree, and extend downward to either leaf entries or references to subordinate naming contexts. It can range in size from a single entry to the entire **DIT**.

Figure 2–3 illustrates valid and invalid naming contexts. Notice that the correct ones on the left are contiguous, and the incorrect ones on the right are not.

***Figure 2–3 Valid and Invalid Naming Contexts***

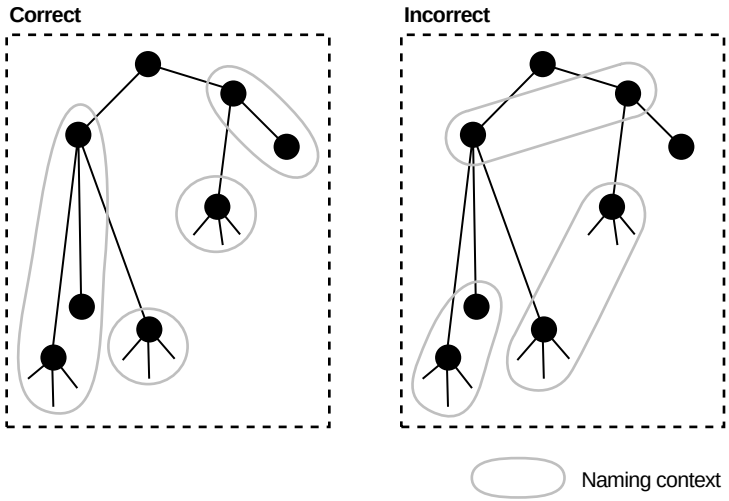


To enable users to search for specific naming contexts, you can publish those naming contexts by using either Oracle Directory Manager or ldapmodify.

> **See Also:** "Managing Naming Contexts" on page 5-16 for instructions on how to publish a naming context

# The Directory Schema

The directory **schema** contains all information about how data is organized in the DIT—that is, metadata such as **object classes**, **attributes**, **matching rules**, and syntaxes. The directory schema stores this information in a special class of entry called a **subentry**. Oracle Internet Directory, following LDAP Version 3 standards, holds schema definitions in the subentry called `subSchemaSubentry`.

You can add new object classes and objects by modifying `subSchemaSubentry`. You cannot, however, add new matching rules and syntaxes beyond those already supported by Oracle Internet Directory.

**See Also:**

- Chapter 6, "Managing the Directory Schema."
- Appendix E for a list of both standard and proprietary schema elements installed with Oracle Internet Directory

# Security

Oracle Internet Directory provides many powerful mechanisms for protecting information from unauthorized access.

This section contains these topics:

- Authentication: Ensuring that the identities of users, hosts, and clients are correctly validated
- Access Control and Authorization: Ensuring that a user reads or updates only the information for which that user has privileges
- Data Integrity: Ensuring that data is not modified during transmission
- Data Privacy: Ensuring that data is not disclosed during transmission
- Password Encryption: Ensuring protection of user passwords through any of four encryption options

## Authentication

Authentication is the process by which the directory server establishes the true identity of the user connecting to the directory. It occurs when an LDAP session is established by means of the ldap-bind operation. Thus every session has an associated user identity. This identity is also referred to as an authorization ID.

To ensure that the identities of users, hosts, and clients are correctly known, Oracle Internet Directory provides three authentication options: anonymous, simple, and SSL.

### Anonymous Authentication

If your directory is available to everyone, then you can allow users to log in to the directory anonymously. When using **anonymous authentication**, users simply leave blank the user name and password fields when they log in. Each anonymous user then exercises whatever privileges are specified for anonymous users.

### Simple Authentication

In this case, the client identifies itself to the server by means of a DN and a password which are not encrypted when sent over the network. In the **simple authentication** option, the server verifies that the DN and password sent by the client match the DN and password stored in the directory.

### Authentication Using Secure Sockets Layer (SSL)

**SSL** is an industry standard protocol for securing network connections. It provides authentication through the exchange of **certificates** that are verified by trusted certificate authorities. A certificate ensures that an entity's identity information is correct. An entity can be an end user, a database, an administrator, a client, or a server. A **certificate authority (CA)** is an application that creates public key certificates that are given a high level of trust by all the parties involved.

You can use SSL in one of three authentication modes:

| SSL Mode | Description |
|---|---|
| No authentication | Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, only SSL encryption / decryption is used. |
| One-way authentication | Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |
| Two-way authentication | Both client and server authenticate themselves to each other. Both the client and server send certificates to each other. |

**Components of SSL** The components of SSL include:

- **Certificate**

  A certificate ensures that the entity's identity information is correct and that the public key actually belongs to that entity. A certificate is created when an entity's public key is signed by a trusted identity, that is, a certificate authority (CA). A certificate contains the entity's name, public key, serial number, and expiration date. It may contain information about the privileges associated with the certificate. Finally, it contains information about the CA that issued it. A certificate is valid until it expires or is revoked.

- **Certificate Authority (CA)**

  A certificate authority is a trusted third party that certifies that other entities are who they say they are. The certificate authority verifies the entity's identity and grants a certificate, signing it with the certificate authority's **private key**.

  Different CAs may have different identification requirements when issuing certificates. For example, one certificate authority may want to see a user's driver's license, another may want the certificate request form to be notarized, yet another may want fingerprints of the person requesting a certificate. The certificate authority publishes its own certificate which includes its public key.

  Each network entity has a list of the certificates of the CAs it trusts. Before communicating with another entity, a given entity uses this list to verify that the signature on the other entity's certificate is from a trusted CA. Network entities can obtain their certificates from the same or from different CAs.

- **Wallet**

  A **wallet** is an abstraction used to store and manage authentication data such as keys, certificates, and **trusted certificates**, also called **trustpoints**, which are needed by SSL. In an Oracle environment, each system using SSL has a wallet with an **X.509** version 3 certificate, private key, and list of trusted certificates.

  Security administrators use the **Oracle Wallet Manager** to manage security credentials on the server. Wallet owners use it to manage security credentials on clients. Specifically, the Oracle Wallet Manager is used to do the following:

  - Generate a **public/private key pair** and create a certificate request for submission to a certificate authority
  - Install a certificate for the entity
  - Configure trusted certificates for the entity

**See Also:** Appendix C for detailed information about managing wallets by using the Oracle Wallet Manager

**The SSL Handshake** At the beginning of their communication, the client and directory server perform a **handshake** which includes three important steps:

- The client and server establish which cipher suite to use. A **cipher suite** is a set of authentication, encryption, and data integrity algorithms used for exchanging messages between network nodes. During an SSL handshake, the two nodes negotiate to see which cipher suite they will use when transmitting messages back and forth.
- The server sends its certificate to the client. The client verifies that the directory server's certificate was signed by a trusted CA.

  Similarly, if client authentication is required, the client sends its own certificate to the directory server. The directory server verifies that the client's certificate was signed by a trusted CA.
- The client and directory server exchange key material using **public-key cryptography**, and, from this material, they each generate a **session key**. All subsequent communication between client and directory server is encrypted and decrypted by using this set of session keys and the negotiated cipher suite.

**See Also:**

- "Data Integrity: Ensuring that data is not modified during transmission" on page 2-12 for more information on data integrity and encryption
- "Supported Cipher Suites" on page 8-2 for a list of SSL cipher suites supported in Oracle Internet Directory

**SSL and Oracle Internet Directory** SSL authentication between a client and a directory server involves three basic steps:

1. The user initiates an LDAP connection to the directory server by using SSL on the SSL port. (The default SSL port is 636.)
2. SSL performs the handshake between client and directory server.
3. If the handshake is successful, the directory server verifies that the user has the appropriate authorization to access the directory.

**See Also:**

- Chapter 8, "Managing Secure Sockets Layer (SSL)"
- Chapter 9, "Managing Directory Access Control" for instructions on setting access control policies
- Appendix C, "Using Oracle Wallet Manager" for a discussion of certificates and wallets

## Access Control and Authorization

Authorization is the process of ensuring that a user reads or updates only the information for which that user has privileges. When directory operations are attempted within a directory session, the directory server ensures that the user—identified by the authorization ID associated with the session—has the requisite permissions to perform those operations. Otherwise, the operation is disallowed. Through this mechanism, the directory server protects directory data from unauthorized operations by directory users. This mechanism is called access control.

Access control information is the directory metadata that captures the administrative policies relating to access control.

ACI is stored in Oracle Internet Directory as user-modifiable operational attributes. Typically, a list of these ACI attribute values, called an Access Control List (ACL), is associated with directory objects. The attribute values on that list govern the access policies for those directory objects.

ACIs are represented and stored as text strings in the directory. These strings must conform to a well defined format. Each valid value of an ACI attribute represents a distinct access control policy. These individual policy components are referred to as ACI Directives or **ACIs** and their format is called the ACI Directive format.

Access control policies can be prescriptive, that is, their security directives can be set to apply downward to all entries at lower positions in the **directory information tree (DIT)**. The points from which such access control policies apply are called **Access Control Policy Points** (**ACPs**).

**See Also:**

- Chapter 9, "Managing Directory Access Control" for instructions on setting access control policies
- Appendix D, "Using Access Control Directive Format" for instructions on correctly formatting ACI directives

## Data Integrity

Oracle Internet Directory ensures that data has not been modified, deleted, or replayed during transmission by using SSL. This SSL feature generates a cryptographically secure message digest—through cryptographic checksums using either the **MD5** algorithm or the **Secure Hash Algorithm (SHA)**—and includes it with each packet sent across the network.

## Data Privacy

Oracle Internet Directory ensures that data is not disclosed during transmission by using **public-key encryption** available with Secure Sockets Layer (SSL). In public-key encryption, the sender of a message encrypts the message with the public key of the recipient. Upon delivery, the recipient decrypts the message using the recipient's private key. Specifically, Oracle Internet Directory supports two levels of encryption available through SSL:

- DES40

  The DES40 algorithm, available internationally, is a variant of **DES** in which the secret key is preprocessed to provide forty effective **key** bits. It is designed for use by customers outside the USA and Canada who want to use a DES-based encryption algorithm. This feature gives commercial customers a choice in the algorithm they use, regardless of their geographic location.

- RC4_40

  Oracle has obtained license to export the RC4 data encryption algorithm with a 40-bit key size to virtually all destinations where other Oracle products are available. This makes it possible for international corporations to safeguard their entire operations with fast cryptography.

  **See Also:** Chapter 8, "Managing Secure Sockets Layer (SSL)" for more information about SSL

## Password Encryption

During installation, you were prompted to set the encryption scheme for passwords. You can change that initial configuration by using either Oracle Directory Manager or ldapmodify. You must be a superuser to change the type of password encryption.

To encrypt passwords, Oracle Internet Directory uses the MD4 algorithm as the default. MD4 is a one-way hash function that produces a 128-bit hash, or message digest. You can change this default to one of the following:

- No Encryption
- MD5—An improved, and more complex, version of MD4
- SHA—Secure Hash Algorithm, which produces a 160-bit hash, longer than MD5. The algorithm is slightly slower than MD5, but the larger message digest makes it more secure against brute-force collision and inversion attacks.
- UNIX Crypt—The UNIX encryption algorithm

The value you specify is stored in the `orclCryptoScheme` attribute in the **Root DSE**. This attribute is single-valued.

During authentication to a directory server, a user enters a password in clear text. The server then hashes the password by using the specified encryption algorithm, and verifies it against the hashed password in the `userPassword` attribute. If the hashed password values match, then the server authenticates the user. If the hashed password values do not match, then the server sends the user an Invalid Credentials error message.

**See Also:** "Managing Password Encryption" on page 5-17

## National Language Support

Oracle Internet Directory follows LDAP Version 3 internationalization (I18N) standards. These standards require that the database storing directory data use the **UTF-8** (Unicode Transformation Format 8-bit) character set. This allows Oracle Internet Directory to store the character data of almost any language that Oracle NLS supports. Moreover, although several different **Application Program Interfaces (APIs)** are involved in the Oracle Internet Directory implementation, Oracle Internet Directory ensures that the correct character encoding is used with each API.

NLS uses both single-byte and multi-byte characters. A single-byte character is represented by one byte of memory. ASCII text, for example, uses single-byte characters. By contrast, a multi-byte character can be represented by more than one byte. Simplified Chinese, for example, uses multi-byte characters. A directory entry in simplified Chinese might look like this:

```
dn: o=\274\327\271\307\316\304,c=\303\300\271\372
objectclass: top
objectclass: organization
o: \274\327\271\307\316\304
```

where the attribute values correspond to character strings in the simplified Chinese character set.

The main Oracle Internet Directory components—OID Monitor (OIDMON), OID Control Utility (OIDCTL), Oracle directory server (OIDLDAPD), and Oracle directory replication server (OIDREPLD)—always use the UTF-8 character set by default.

Oracle Directory Manager, a Java-based tool, internally uses Unicode (**UCS2**—that is, fixed-width 16-bit **Unicode**). In Java, UCS2 is the easiest way to handle characters—including English characters. The Java client uses standard Java packages to convert both to and from UCS2 and UTF-8. This enables Oracle Directory Manager to handle the LDAP Version 3 protocol using UTF-8.

**See Also:**

- "Oracle Internet Directory Architecture" on page 2-20 for information on the main Oracle Internet Directory components
- Chapter 12, "Managing National Language Support (NLS)" for instructions on using NLS in Oracle Internet Directory
- *Oracle8i National Language Support Guide* for a detailed discussion of NLS

# Oracle Internet Directory Architecture

This section contains these topics:



## An Oracle Internet Directory Node

Figure 2–4 shows the various directory server components and their relationships running on a single node.

Net8 is used for all connections between the Oracle database server and:

- The **OID Control Utility**
- The LDAP Server Instance 1 non-SSL port 389
- The LDAP Server Instance 2 SSL Enabled port 636
- The **OID Monitor**

LDAP is used for connections between The LDAP Server Instance 1 non-SSL port 389 and:

- Oracle Directory Manager
- Oracle directory replication server

The two LDAP server instances and the replication server connect to OID Monitor by way of the operating system.

*Figure 2–4 A Typical Oracle Internet Directory Node*

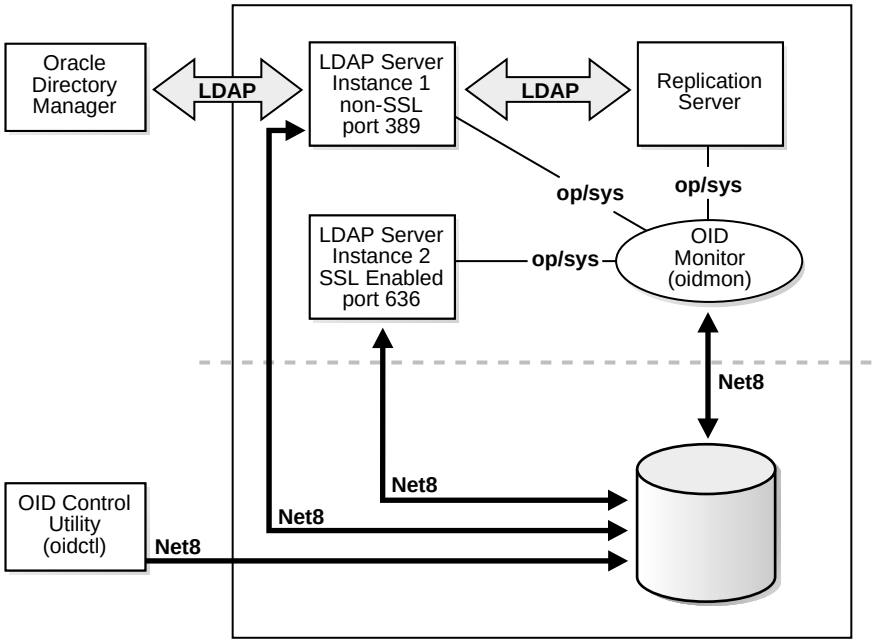


**Note:** In Figure 2–4, the database is on the same node as the directory server processes. However, because all connections with the database are through **Oracle Call Interface (OCI)** and **Net8**, it is possible to use a database on a different server.

An Oracle Internet Directory node (Figure 2–4) includes the following major components:

| Component | Description |
|---|---|
| LDAP server instance | Also called an Oracle directory server instance. It services directory requests through a single Oracle Internet Directory dispatcher process listening at a specific TCP/IP port number. There can be more than one LDAP server instance on a node, each listening on a different port. Oracle Internet Directory dispatcher and server processes use multiple threads. |
| Replication server | Also called an Oracle directory replication server. It tracks and sends changes to replicated servers in an Oracle Internet Directory system. There can be only one replication server on a node. You can choose whether or not to install and use the replication server. |
| Oracle8*i* database | Stores the directory data. Oracle Corporation strongly recommends that you dedicate a database for use by the directory. The database can reside on the same node as the servers, or on a separate node. |
| **OID Monitor** (OIDMON) | Initiates, monitors, and terminates the LDAP server processes. If you elect to install a replication server, OID Monitor controls it. When you issue commands through OID Control Utility (OIDCTL) to start or stop directory server instances, your commands are interpreted by this process.<br><br>OID Monitor executes the LDAP server instance startup and shutdown requests that you initiate from OID Control Utility. OID Monitor also monitors servers and restarts them if they have stopped running for abnormal reasons.<br><br>When it starts a server instance, OID Monitor adds an entry into the directory instance registry and updates data in a process table. When it shuts down the directory server instance, it deletes the registry entry as well as the data corresponding to that particular instance from the process table. If OID Monitor restarts a server that has stopped abnormally, it updates the registry entry with the start time of the server.<br><br>All OID Monitor activity is logged in the file `ORACLE_HOME/ldap/log/oidmon.log`. This file is on the Oracle Internet Directory server file system.<br><br>OID Monitor checks the state of the servers through mechanisms provided by the operating system. |

| Component | Description |
|---|---|
| OID Control Utility (OIDCTL) | Communicates with OID Monitor by placing message data in Oracle Internet Directory server tables. This message data includes configuration parameters required to run each Oracle directory server instance. |

The Oracle directory replication server uses LDAP to communicate with an Oracle directory (LDAP) server instance. To communicate with the database, all components use OCI / Net8. Oracle Directory Manager and the command line tools communicate with the Oracle directory (LDAP) servers over LDAP.

## An Oracle Directory (LDAP) Server Instance

Each Oracle directory (LDAP) server instance looks similar to what Figure 2–5 illustrates.

***Figure 2–5 LDAP Server Instance Architecture***

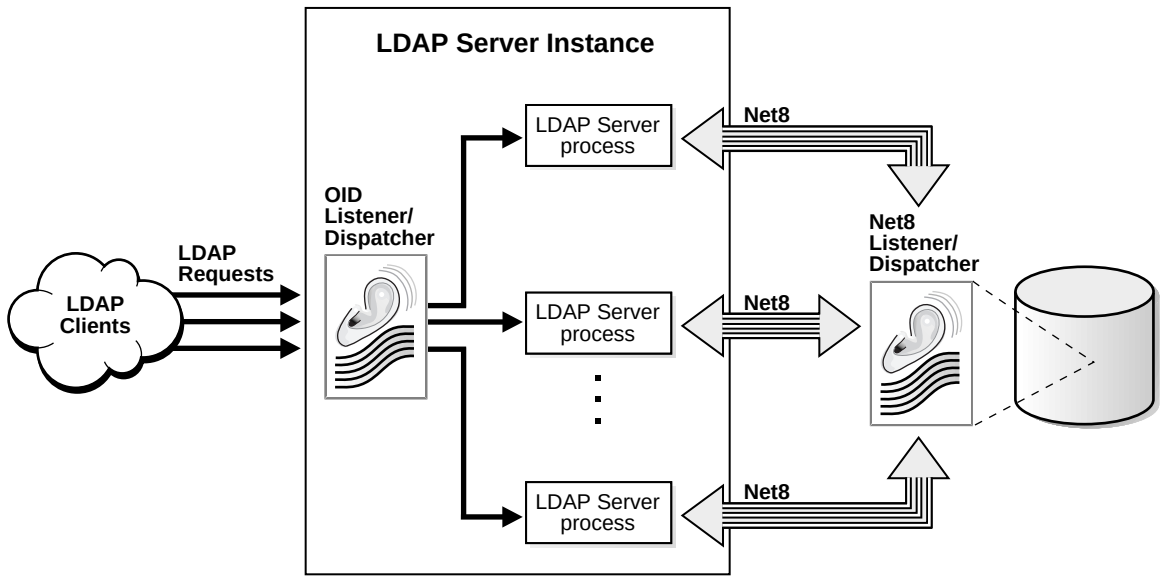


LDAP clients send LDAP requests to an Oracle Internet Directory listener / dispatcher process listening for LDAP commands at its port.

The OID listener / dispatcher launches the LDAP directory server which, in turn creates server processes. Multiple server processes allow Oracle Internet Directory to take advantage of multiple processor systems. The number of server processes created is determined by a configuration parameter (ORCLSERVERPROCS). The default is 1 (one). A worker thread for each operation processes the client request.

Database connections from each server process are spawned as needed, up to a maximum number determined by a configuration parameter (ORCLMAXCC). The default value for this parameter is 10. The server processes communicate with the data server by way of Net8. A Net8 Listener / Dispatcher relays the request to the Oracle data server.

## Configuration Set Entries

The configuration parameters for each Oracle directory server instance are stored in a directory entry called a configuration set entry, or configset. A configuration set entry holds the configuration parameters for a specific instance of the directory server. When you start an instance of a server using OID Control Utility, the command contains a reference to one of these configsets and uses the information it contains.

The Oracle directory server is installed with a default configuration set entry (configset0) so that you can run the directory server immediately. You can create customized configuration set entries by adding new ones that change specific parameters to meet your needs. You can view, add, and modify these entries by using either **Oracle Directory Manager** or the appropriate command line tool.

**See Also:**

- "Managing Server Configuration Set Entries" on page 5-2
- "Configuration Set Entry Attributes" on page E-5 for a list of configuration set entry attributes

## Example: How Oracle Internet Directory Works

Now that all of the concepts are introduced, the following example shows you how Oracle Internet Directory processes a search request.

1. The user or client enters a search request that is conditioned by one or more of the following options:
   - SSL: The client and server can establish a session that uses SSL encryption and authentication, or SSL encryption only. If SSL is not used, the client's message is sent in clear text.
   - Type of user: The user can seek access to the directory either as a particular user or as an anonymous user, depending on which of the two has the necessary privileges to perform the desired function.
   - Filters: The user can narrow the search by using one or more search filters, including those that use the Boolean conditions "and," "or," and "not," and those that use other operators such as "greater than, "equal to," and "less than".
2. If the user or client issues the command by using Oracle Directory Manager, the latter invokes a query function in the Java Native Interface which, in turn, invokes a function in the C API. If the user or client uses a command line tool, then the tool directly invokes a C function in the C API.
3. The C API, using the LDAP protocol, sends a request to a directory server instance to connect to the directory.
4. The directory server authenticates the user, a process called binding. The directory server also checks the Access Control Lists (ACLs) to verify that the user is authorized to perform the requested search.
5. The directory server converts the search request from LDAP to Oracle Call Interface (OCI) / Net8 and sends it to the Oracle8*i* database.
6. The Oracle8*i* database retrieves the information and passes it back through the chain—to the directory server, then to the C API, and, finally, to the client.

# Distributed Directories: An Overview

Although an online directory is logically centralized, it can physically distribute its data onto several servers. Physical distribution reduces the work a single server would otherwise have to do. It also enables the directory to accommodate a larger number of entries.

A distributed directory can be either replicated or partitioned. When information is replicated, the same naming contexts are stored by more than one server. When information is partitioned, each directory server stores one or more unique, non-overlapping naming contexts. In a distributed directory, some information may be partitioned and some may be replicated.

**See Also:**

- "Distributed Directories: Replication" on page 2-26
- "Distributed Directories: Partitioning" on page 2-42

# Distributed Directories: Replication

Replication is one way of distributing directory information. It improves performance by providing more servers to handle queries. It improves reliability by eliminating risks associated with a single point of failure.

Each copy of a naming context that is contained within a server is called a replica. A directory server can hold both read-only and updatable replicas.

Servers that hold updatable replicas are called suppliers. Their changes are propagated to other servers called consumers.

You specify how many times the replication server should try to apply changes to consumers. Once that number is reached, the replication server moves the changes to a human intervention queue, then attempts to apply them at regular, less frequent intervals that you specify.

Figure 2–6 shows a replicated directory.

***Figure 2–6 A Replicated Directory***

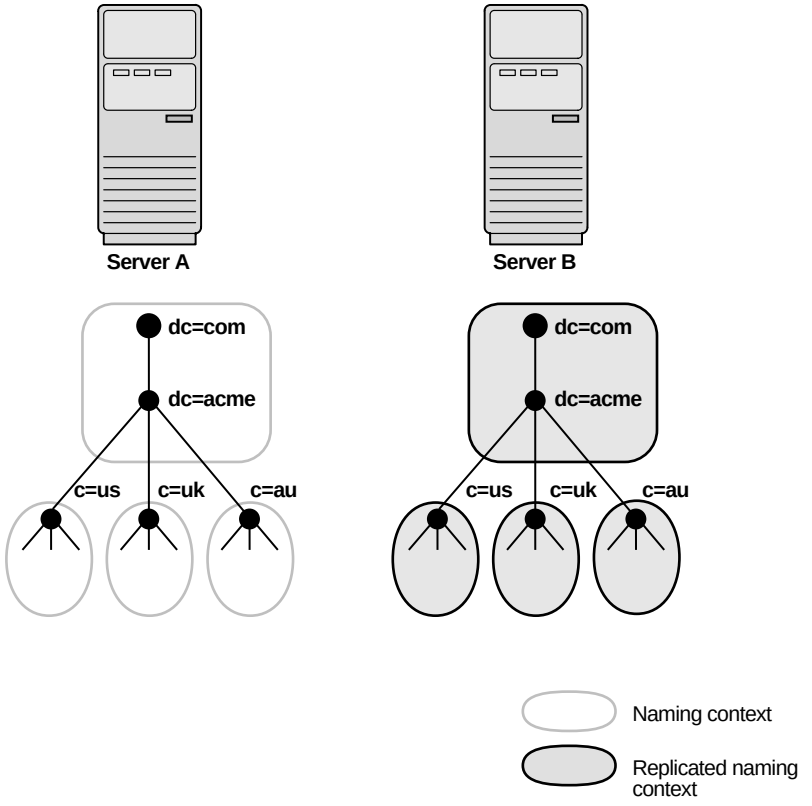


**Note:** This release of Oracle Internet Directory enables replication at the level of the naming context. It does not support replication of part of a naming context.

Also, although there are no Internet standards for directory replication yet, such standards are being developed by the IETF. Oracle Internet Directory replication adheres to the IETF standard proposal for representing directory change information in **change logs**.

> **See Also:** "Replication Architecture" on page 2-30 for more on change logs

This section contains these topics:

- Directory Replication Groups and Replication Agreements
- Oracle Advanced Symmetric Replication (ASR)
- Replication Architecture
- Change Log Purging
- Conflict Resolution in Replication
- How Replication Works: An Overview
- How Replication Works: A Closer Look

> **See Also:** Chapter 10, "Managing Directory Replication" for more information on replication

## Directory Replication Groups and Replication Agreements

The set of directory servers that participate in replication of a given naming context is called a Directory Replication Group (DRG). A special directory entry, called a replication agreement, represents the replication relationship among the directory servers in a DRG.

It is possible for a directory server to be both a supplier and a consumer of change log information. Oracle Internet Directory uses this feature to support multimaster replication.

Figure 2–7 illustrates a Directory Replication Group in which three nodes share updates with each other in a replication agreement.

***Figure 2–7 Directory Replication Group***

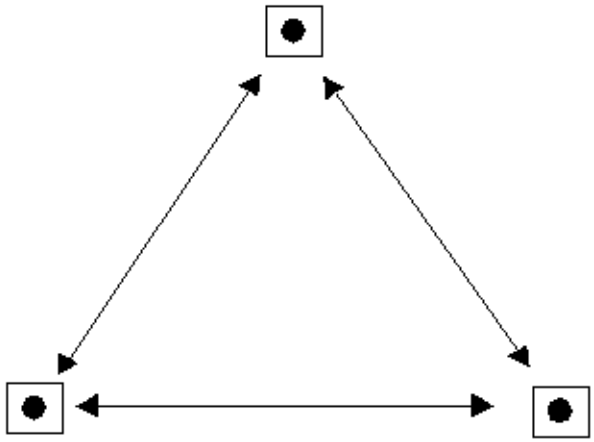

In Figure 2–7, each bullet represents a node of Oracle Internet Directory. The agreement is identical on each node except for local options such as partitioned naming contexts on the local directory server. The replication agreement on each node lists all the other nodes to which it delivers, and from which it receives, changes.

> **See Also:** "Task 5: Configure Replication" on page 10-9 for information about how to configure replication agreements

## Oracle Advanced Symmetric Replication (ASR)

Transport of update information between nodes in a replication agreement is managed by Oracle Advanced Symmetric Replication (ASR), a store-and-forward transport feature available in Oracle8*i*. It allows database tables to be kept synchronized across two Oracle databases.

ASR stores local changes and periodically propagates them in batches to consumer servers. The consumer replication servers apply the remote changes to the local directory server and then purge the applied remote changes from their local stores.

ASR environments allow read and update access to directory tables anywhere in the system. Typical ASR configurations use row-level replication with asynchronous data propagation.

ASR provides proven network tolerance and its data transfer can be controlled and monitored by Oracle Enterprise Manager. Such manageability allows a high degree of flexibility in how the data transfer is scheduled.

**See Also:** *Oracle8i Replication* for information about ASR

## Replication Architecture

Supplier servers write their changes to change logs, and then regularly send batched directory changes to other supplier and consumer servers. Consumer servers receive the change log data, then reproduce the changes locally.

When you configure replication, you specify which nodes in a replication group share changes. Regardless of the number of nodes you introduce into the replication environment, the basic architecture for replication remains the same. Local changes are distributed to remote nodes and applied by replication server processing. To apply the changes on a remote node, the replication server, acting as a client, sends commands to the directory server that implements them.

**See Also:** "Task 5: Configure Replication" on page 10-9 for information on configuring replication

## Change Log Purging

Change log purging takes place in Oracle Internet Directory in two ways:

| | |
|---|---|
| Change number-based | This is the default method. The replication server purges those changes that have already been applied to all the nodes in a DRG. |
| Time-based | You can run this method to augment change number-based purging. To use this additional method, you set a parameter specifying in hours the lifespan of change log objects. For example, you can set this parameter to purge all change log objects that are 24 hours old. Use this method to prevent the change log from becoming too large. |

**See Also:**

- "Oracle Directory Replication Server Parameters" on page 10-10
- "Viewing and Modifying Replication Configuration Parameters by Using Oracle Directory Manager" on page 10-11
- "Modifying Replication Configuration Parameters by Using Command Line Tools" on page 10-12

## Conflict Resolution in Replication

Multimaster replication enables updates to multiple directory servers. Conflicts occur whenever the Oracle directory replication server attempts to apply remote changes from a supplier to a consumer and fails for some reason.

Four kinds of LDAP operations can lead to conflicts:

- Addition
- Deletion
- Modification
- Modification of either an RDN or a DN

### Levels at Which Replication Conflicts Occur

There are two types of conflicts:

- Entry-level conflicts
- Attribute-level conflicts

**Entry-Level Conflicts** An entry-level conflicts is caused when the Oracle directory replication server attempts to apply a change to the consumer. Such a change could be one of the following types of changes to the consumer:

- Adding an entry that already exists
- Deleting an entry that does not exist
- Modifying an entry that does not exist
- Applying a modifyrdn operation when the DN does not exist

These conflicts can be difficult to resolve. For instance, it may be impossible to resolve a conflict because:

- The entry has been moved to a different location
- The entry has not yet arrived from a supplier
- The entry has been deleted
- The entry never existed

If an entry exists and it should not, then that may mean:

- The entry was added earlier
- The entry recently underwent a modifydn operation

**Attribute-Level Conflicts** An attribute-level conflict is caused when two directories are updating the same attribute with different values at different times. If the attribute is single-valued, then the replication process resolves the conflict by examining the timestamps of the changes involved in the conflict.

### Typical Causes of Conflicts

Conflicts usually stem from the timing of changes arising from the occasional slowness or transmission failure over wide-area networks. Also, an earlier inconsistency might continue to cause conflicts if it is not resolved in a timely manner.

### Automated Resolution of Conflicts

The Oracle directory replication server attempts to resolve all conflicts that it encounters by following this process:

1. The conflict is detected when a change is applied.
2. The replication process attempts to reapply the change a specific number of times or repetitively for a specific amount of time after a specific waiting period.
3. If the replication process reaches the retry limit without successfully applying the change, then it flags the change as a conflict and moves the change to a low-priority, human intervention queue. Changes are then applied according to the time unit specified in the `orclHIQSchedule` parameter in the replication agreement. Before it moves the change, the Oracle directory replication server writes the conflict into a log file for the system administrator.

**Note:** There is no conflict resolution of schema, catalog, and group entries during replication. This is because attempting resolution of such large multi-valued attributes would have a significant negative impact on performance. Be careful to avoid updating such entries from more than one master at a time.

## How Replication Works: An Overview

Figure 2–8 provides a general overview of the replication process on both the supplier and the consumer. It illustrates the following:

**On the supplier side:**

1. When an LDAP client issues a directory modification, the LDAP server generates a change log object in the change log object store.
2. At a scheduled time, the replication server spawns an outbound change log processing thread to translate the change log object into a row—for example, Change entry—in the change log table.
3. When a change entry gets inserted and then committed to the change log table, ASR immediately copies the change into the deferred transaction queue.
4. After a scheduled interval, ASR pushes pending transactions from the deferred transaction queue across the network to the consumer change log table.

**On the consumer side:**

1. The replication server spawns a change log processing thread for each supplier, based on a scheduled replication cycle.
2. The change log processing thread consults the change status table for the last change applied from the supplier to the consumer.
3. It then fetches and applies all the new changes from the change log table to the LDAP server.
4. After completing the change log processing, the change log processing thread updates the change status table to record the last change applied from the supplier before exiting.
5. ASR copies the change status update into the deferred transaction queue.
6. After the scheduled ASR replication interval, ASR pushes pending change status updates from the deferred transaction queue to the supplier change status table.

Figure 2–8 illustrates the replication process.

***Figure 2–8 Overview of ASR-Based Replication Architecture***

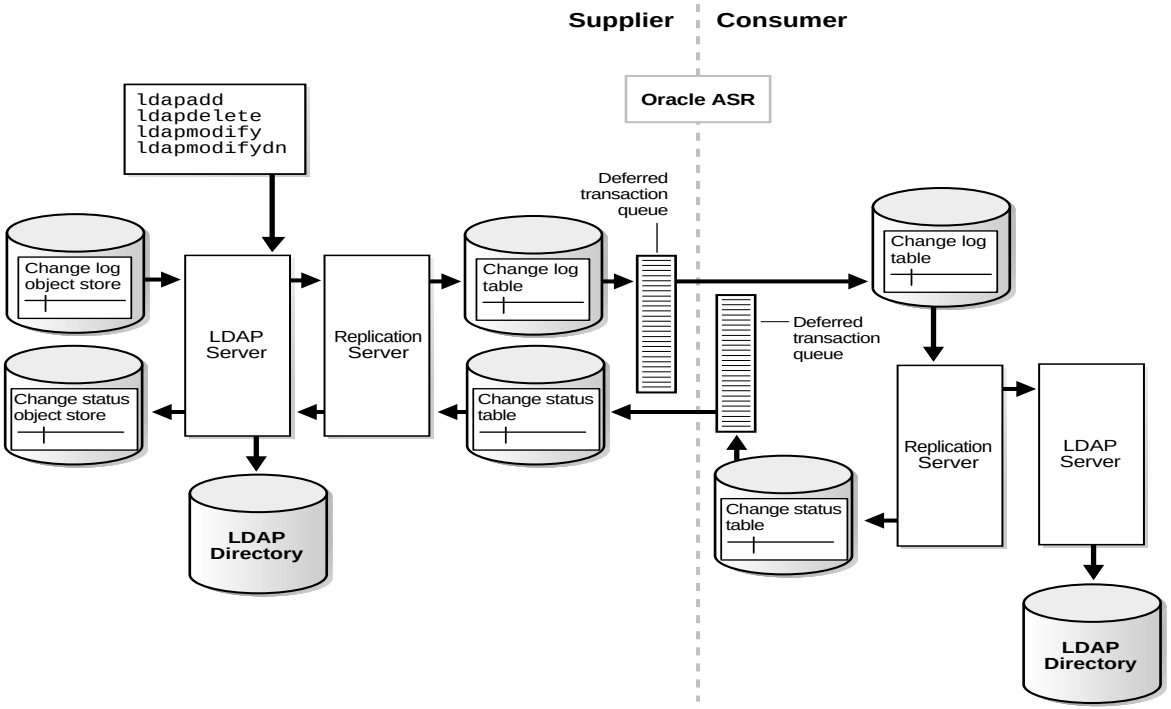


Although, in Figure 2–8, the roles of supplier and consumer have been separated, in an actual multimaster replication environment, each directory server is both a supplier and a consumer. In such an environment, the purging of entries that are already applied or that have been dropped as candidate changes occurs regularly. Remote change records in the local Changelog table are purged by the garbage collection thread if they have been applied locally. Local change records in the local Changelog table are purged by the garbage collection thread if they have been distributed to all the consumers.

**See Also:**

- "How Replication Works: An Overview" on page 2-33 for a more detailed explanation of how the replication server adds, deletes, and modifies entries, as well as how it modifies DNs and RDNs
- "Conflict Resolution in Replication" on page 2-31 for a detailed description of how the replication server resolves conflicts when it adds, deletes, and modifies entries and when it modifies DNs and RDNs

## How Replication Works: A Closer Look

This section describes in more detail than the previous section how the automated replication process adds, deletes, and modifies entries, and how it modifies DNs and RDNs.

### How the Replication Process Adds a New Entry to a Consumer

When Oracle directory replication server successfully adds a new entry to a consumer, it follows this change application process:

1. The Oracle directory replication server looks in the consumer for the DN of the parent of the target entry. Specifically, it does this by looking for a **global unique identifier (GUID)** assigned to the DN of the parent.
2. If the parent entry exists, then the Oracle directory replication server composes a DN for the new entry and places the new entry under its parent in the consumer. It then places the change entry in the purge queue.

**If the change entry is not successfully applied on the first try:**

The Oracle directory replication server places the new change entry in the retry queue, sets the number of retries to the configured maximum, and repeats the change application process.

**If the change entry is not successfully applied on *all but the last* retry:**

The Oracle directory replication server keeps the change entry in the retry queue, decrements the number of retries, and repeats the change application process.

**If the change entry is not successfully applied on the last retry:**

The Oracle directory replication server checks to see if the new entry is a duplicate of an existing entry.

**If the change entry is a duplicate entry**:

The Oracle directory replication server applies the following conflict resolution rules:

* The entry with the older creation time stamp is used.
* If both entries have the same creation time stamp, then the entry with the smaller GUID is used.

If the change entry is used, then the target entry is removed, the change is applied, and the change entry is placed in the purge queue.

If the target entry is used, then the change entry is placed in the purge queue.

**If the change entry is not a duplicate entry:**

The Oracle directory replication server places the change entry in the human intervention queue, and repeats the change application process at the interval you specified in the `orclHIQSchedule` parameter.

**If the change entry is not successfully applied after it has been placed in the human intervention queue:**

The Oracle directory replication server keeps the change in this queue, and repeats the change application process at specified intervals while awaiting action by the administrator. The administrator can use the OID reconciliation tool and the human intervention queue manipulation tool to resolve the conflict.

### How the Replication Process Deletes an Entry

When the Oracle directory replication server deletes an entry from a consumer, it follows this change application process:

1. The Oracle directory replication server looks in the consumer for an entry with a GUID matching the one in the change entry.
2. If the matching entry exists in the consumer, then the Oracle directory replication server deletes it. It then places the change entry in the purge queue.

**If the change entry is not successfully applied on the first try:**

The Oracle directory replication server places the change entry in the retry queue, sets the number of retries to the configured maximum, and repeats the change application process.

**If the change entry is not successfully applied on *all but the last* retry:**

The Oracle directory replication server keeps the change entry in the retry queue, decrements the number of retries, and repeats the change application process.

**If the change entry is not successfully applied on the last retry:**

The Oracle directory replication server places the change entry in the human intervention queue and repeats the change application process at specified intervals.

**If the change entry is not successfully applied after it has been placed in the human intervention queue:**

The Oracle directory replication server keeps the change entry in this queue, and repeats the change application process at specified intervals while awaiting action by the administrator. The administrator can use the OID reconciliation tool and the human intervention queue manipulation tool to resolve the conflict.

### How the Replication Process Modifies an Entry

When the Oracle directory replication server modifies an entry in a consumer, it follows this change application process:

1. The Oracle directory replication server looks in the consumer for an entry with a GUID matching the one in the change entry.
2. If the matching entry exists in the consumer, then the Oracle directory replication server compares each attribute in the change entry with each attribute in the target entry.
3. The Oracle directory replication server then applies the following conflict resolution rules:
   a. The entry with the most recent modify time is used.
   b. The entry with the most recent version of the attribute is used.
   c. The modified attribute on the host whose name is closest to the beginning of the alphabet is used.
4. The Oracle directory replication server applies the filtered modification, and places the change entry in the purge queue.

**If the change entry is not successfully applied on the first try:**

The Oracle directory replication server places the change entry in the retry queue, sets the number of retries to the configured maximum, and repeats the change application process.

**If the change entry is not successfully applied on *all but the last* retry:**

The Oracle directory replication server keeps the change entry in the retry queue, decrements the number of retries, and repeats the change application process.

**If the change entry is *not* successfully applied by the last retry:**

The Oracle directory replication server places the change entry in the human intervention queue and repeats the change application process at specified intervals.

**If the change entry is not successfully applied after it has been placed in the human intervention queue:**

The Oracle directory replication server keeps the change entry in this queue, and repeats the change application process at specified intervals while awaiting action by the administrator. The administrator can use the OID reconciliation tool and the human intervention queue manipulation tool to resolve the conflict.

### How the Replication Process Modifies a Relative Distinguished Name

When the Oracle directory replication server modifies the RDN of an entry in a consumer, it follows this change application process:

1. The Oracle directory replication server looks in the consumer for the DN with a GUID that matches the GUID in the change entry.
2. If the matching entry exists in the consumer, then the Oracle directory replication server modifies the RDN of that entry and places the change entry in the purge queue.

**If the change entry is not successfully applied on the first try:**

The Oracle directory replication server places the change entry in the retry queue, sets the number of retries to the configured maximum, and repeats the change application process.

**If the change entry is not successfully applied on *all but the last* retry:**

The Oracle directory replication server keeps the change entry in the retry queue, decrements the number of retries, and repeats the change application process.

**If the change entry is not successfully applied on the last retry:**

The Oracle directory replication server places the change entry in the human intervention queue and checks to see if it is a duplicate of the target entry.

**If the change entry is a duplicate entry**:

The Oracle directory replication server applies the following conflict resolution rules:

* The entry with the older creation time stamp is used.
* If both entries have the same creation time stamp, then the entry with the smaller GUID is used.

If the change entry is used, then the target entry is removed, the change entry is applied, and then placed in the purge queue.

If the target entry is used, then the change entry is placed in the purge queue.

**If the change entry is not a duplicate entry:**

The Oracle directory replication server places the change entry in the human intervention queue, and repeats the change application process at specified intervals.

**If the change entry is not successfully applied after it has been placed in the human intervention queue**:

The Oracle directory replication server keeps the change entry in this queue, and repeats the change application process at specified intervals while awaiting action by the administrator. The administrator can use the OID reconciliation tool and the human intervention queue manipulation tool to resolve the conflict.

### How the Replication Process Modifies a Distinguished Name

When the Oracle directory replication server modifies the DN of an entry in a consumer, it follows this change application process:

1. The Oracle directory replication server looks in the consumer for the DN with a GUID that matches the GUID in the change entry.

   The Oracle directory replication server also looks in the consumer for the parent DN with a GUID that matches the GUID of the new parent specified in the change entry.

2. If both the DN and the parent DN of the target entry exist in the consumer, then the Oracle directory replication server modifies the DN of that entry and places the change entry in the purge queue.

**If the change entry is not successfully applied on the first try:**

The Oracle directory replication server places the change entry in the retry queue, sets the number of retries to the configured maximum, and repeats the change application process.

**If the change entry is not successfully applied on *all but the last* retry:**

The Oracle directory replication server keeps the change entry in the retry queue, decrements the number of retries, and repeats the change application process.

**If the change entry is *not* successfully applied by the last retry:**

The Oracle directory replication server places the change entry in the human intervention queue and checks to see if it is a duplicate of the target entry.

**If the change entry is a duplicate entry**:

The Oracle directory replication server applies the following conflict resolution rules:

* The entry with the older creation time stamp is used.
* If both entries have the same creation time stamp, then the entry with the smaller GUID is used.

If the change entry is used, then the target entry is removed, the change entry is applied, and then placed in the purge queue.

If the target entry is used, then the change entry is placed in the purge queue.

**If the change entry is not a duplicate entry:**

The Oracle directory replication server places the change entry in the human intervention queue, and repeats the change application process at specified intervals.

**If the change entry is not successfully applied after it has been placed in the human intervention queue:**

The Oracle directory replication server keeps the change entry in this queue, and repeats the change application process at specified intervals while awaiting action by the administrator. The administrator can use the OID reconciliation tool and the human intervention queue manipulation tool to resolve the conflict.

# Distributed Directories: Partitioning

Partitioning is another way of distributing directory information. Figure 2–9 shows a partitioned directory in which some naming contexts reside on different servers.

***Figure 2–9 A Partitioned Directory***

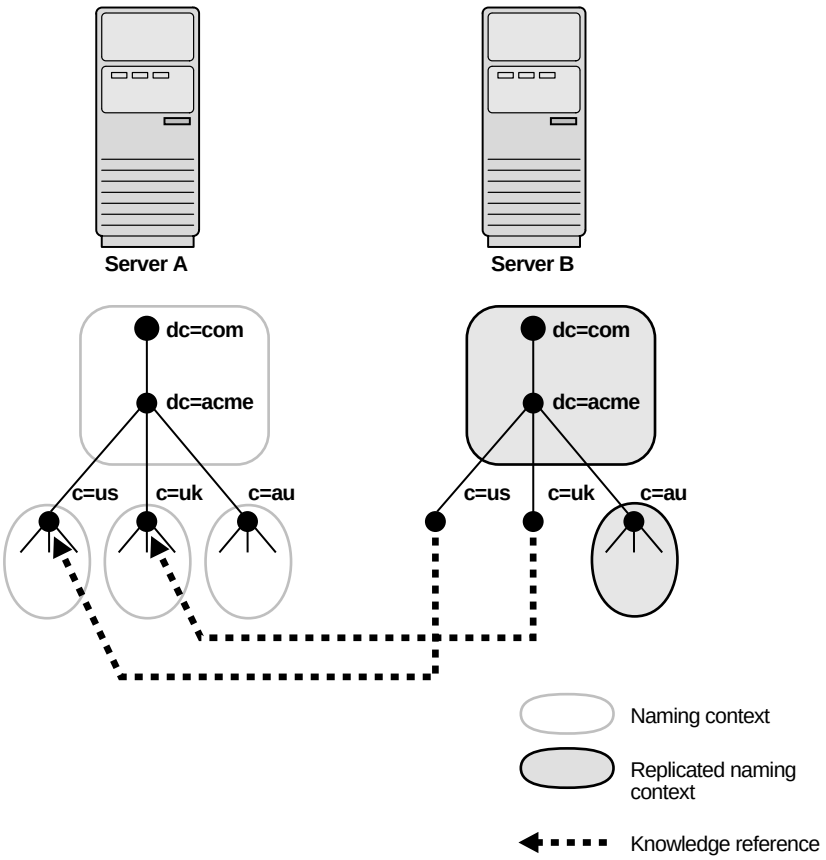


In Figure 2–9, four naming contexts reside on Server A:

- `dc=acme,dc=com`
- `c=us`
- `c=uk`
- `c=au`

Two naming contexts on Server A are replicated on Server B:

- `dc=acme,dc=com`
- `c=au`

The directory uses **knowledge reference**s, also called **referral**s, to locate information that is requested of Server B, but that resides on Server A.

## About Knowledge References (Referrals)

Knowledge references provide the names and addresses of the various naming contexts. In Figure 2–9, Server B uses knowledge references to tell clients that Server A has the requested information in the c=us and c=uk naming contexts. Clients can then use the referral information to contact Server A.

Typically, each directory server contains both superior and subordinate knowledge references. Superior knowledge references point upward in the DIT toward the root. They tie the partitioned naming context to its parent. Subordinate knowledge references point downward in the DIT to other partitions.

For example, in Figure 2–10, Server B holds two naming contexts, each of which is superior to other naming contexts. These two superior naming contexts use subordinate knowledge references to point to their subordinate naming contexts. Conversely, the naming context on Server A has an immediate superior residing on Server B. Server A therefore uses a superior knowledge reference to point to its parent on Server B.

***Figure 2–10 Using Knowledge References to Point to Naming Contexts***

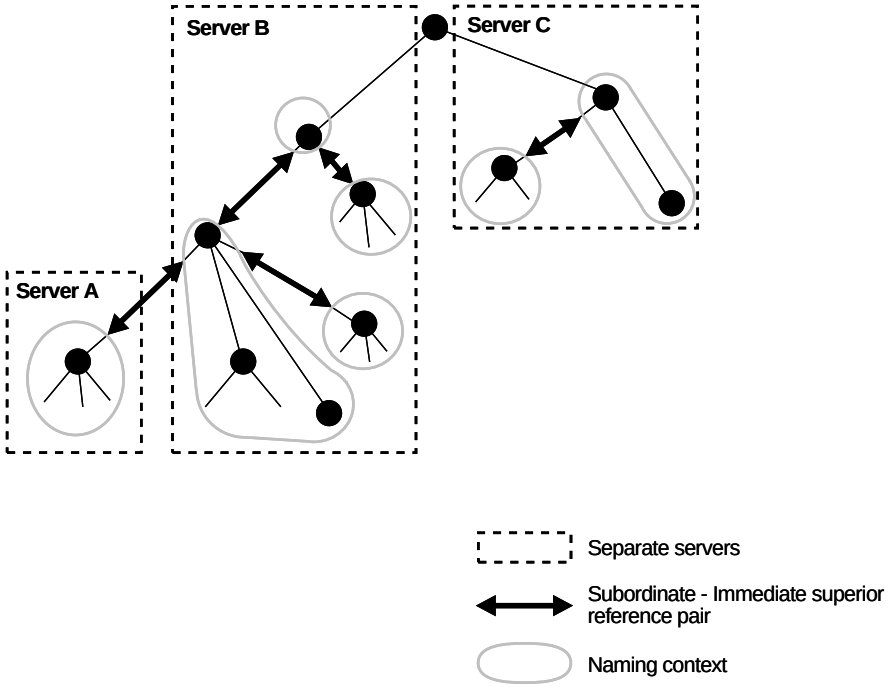


Naming contexts that start at the top of the DIT obviously cannot have a knowledge reference to a superior naming context.

---

**Note:** There are presently no Internet standards for enforcing the validity of knowledge references, and Oracle Internet Directory does not do so. It is up to the administrator to ensure consistency among knowledge references within an enterprise network.

Oracle Corporation recommends that permission for managing knowledge reference entries be restricted like any other privileged administrative function such as schema or access control.

---

## Kinds of Knowledge Reference

There are two kinds of knowledge reference:

**Smart knowledge reference**

Returned when the knowledge reference entry is in the scope of the search. It points the user to the server that stores the requested information.

For example, suppose that:

- Server A holds the naming context `ou=server development,c=us,o=acme`, and has a knowledge reference to Server B
- Server B holds the naming context `ou=sales,c=us,o=acme`

When a user sends a request to Server A for information in `ou=sales,c=us,o=acme`, Server A provides the user with a knowledge reference pointing to Server B.

**Default knowledge reference**

Returned when the base object is not in the directory, and the operation is performed in a naming context not held locally by the server. A default knowledge reference typically sends the user to a server that has more knowledge about the directory partitioning arrangement.

For example, suppose that Server A holds:

- The naming context `c=us,o=acme`
- A knowledge reference to Server PQR that has more knowledge about the overall directory partitioning arrangement

Now suppose that a client requests information on `c=uk,o=acme`. When Server A finds that it does not have the `c=uk,o=acme` naming context, it points the user to Server PQR. From there, the user can find the server holding the requested naming context.

**See Also:** "Managing Knowledge References (Referrals)" on page 7-18

# Synchronizing with Other Directories in a Metadirectory Environment

In a metadirectory environment, you can synchronize multiple directories with Oracle Internet Directory to form a single virtual directory. This section contains these topics:



## About Metadirectories

A metadirectory enables enterprises to synchronize many directories into one "virtual" directory.

Enterprises today often deploy multiple "directories," such as ERP systems, database applications, messaging systems, and Network Operating Systems (NOS). Managing so many different directories can cause problems:

- Redundancy—The same information is represented in many different places in the enterprise.
- High cost of administration—Administrators must maintain essentially the same information in many different places.
- Inconsistent data—Updated information in one directory is not shared with all the other directories.

A metadirectory addresses these problems by integrating all these enterprise directories into one synchronized directory of directories.

## How Oracle Internet Directory Works with a Metadirectory Solution

Oracle Internet Directory release 2.1.1 is interoperable with supported third party metadirectory solutions. This allows various information repositories to synchronize with Oracle Internet Directory and to form a single virtual directory.

A metadirectory solution incorporating Oracle Internet Directory allows:

- Data to be imported from other information repositories, called connected directories, into Oracle Internet Directory
- Data to be exported from Oracle Internet Directory into connected directories

The import and export of data between Oracle Internet Directory and connected directories are performed by software components called metadirectory agents.

Metadirectory vendors provide these as part of their metadirectory solutions. You can also design them by using a metadirectory vendor's framework.

For example, for each employee in an enterprise, Oracle Internet Directory, interoperating with a metadirectory solution, can build a global directory entry. That entry can contain data from such different sources as Human Resources applications, email services, or NOS databases. Users can then access that entry, knowing that it is up-to-date and synchronized among all connected directories.

Moreover, the synchronization can respect the existing data ownership policies—for example, only the Human Resource Directory may be privileged to change an employee's salary attribute. In this way, the metadirectory solution manages directory data and shares information across different directories in an enterprise.

---

**Note:** Oracle Internet Directory release 2.1.1 is interoperable with the supported version of Siemens DirXMetahub.

Oracle Corporation neither licenses the Siemens DirXMetahub product nor ships it on the CD you receive. To obtain the Siemens DirXMetahub, contact Siemens directly at:
`http://www.usa.siemens.com/`

---

A supported metadirectory solution uses Oracle Internet Directory as its metadirectory store—that is, Oracle Internet Directory is the enterprise directory against which other application-specific connected directories are synchronized.

Integrating other directories with Oracle Internet Directory through a metadirectory solution provides these benefits:

- Consistency, data integrity, and better quality of information for both users and applications
- A single point of access for all directory data through standards-based clients like Web browsers, reduced administrative costs, and ease of administration
- A single point of administration using the Oracle Internet Directory management tool
- The ability to keep all connected directories in the environment up-to-date

Oracle Internet Directory support for a metadirectory solution allows these types of directory synchronization:

- Complete and incremental import from connected directories into Oracle Internet Directory (add, modify, and remove operations)
- Complete and incremental export from Oracle Internet Directory into connected directories (add, modify, and remove operations)

**See Also:** Chapter 11, "Synchronizing with Multiple Directories"

# 3

# Preliminary Tasks

This chapter guides you through some tasks you must perform before configuring and using Oracle Internet Directory. It also discusses upgrading from previous releases of Oracle Internet Directory.

Before you can run the administration tools and begin configuring and using the directory, you need to start OID Monitor and start a directory server instance. You also need to reset the default security configuration.

This section contains these topics:

# Task 1: Start the OID Monitor Daemon

The OID Monitor daemon must be running to process commands to start and stop the server.

This section contains these topics:

- Starting the OID Monitor
- Stopping the OID Monitor

## Starting the OID Monitor

To start the OID Monitor:

1. Set the following environment variable to the appropriate language setting. The default language set at installation is AMERICAN_AMERICA.

   ```
   NLS_LANG=APPROPRIATE_LANGUAGE.UTF8
   ```

2. At the system prompt, type:

   ```
   oidmon [connect=net_service_name] [sleep=seconds] start
   ```

| Argument | Description |
|---|---|
| connect=*net_service_name* | Specifies the net service name of the database to which you want to connect. This is the network service name set in the `tnsnames.ora` file. This argument is optional. |
| sleep=*seconds* | Specifies number of seconds after which the OID Monitor should check for new requests from OID Control and for requests to restart any servers that may have stopped. The default sleep time is 10 seconds. This argument is optional. |
| start | Starts the OID Monitor process |

For example:

```
oidmon connect=dbs1 sleep=10 start
```

### Stopping the OID Monitor

To stop the OID Monitor daemon, at the system prompt, type:

```
oidmon [connect=net_service_name] stop
```

| Argument | Description |
|---|---|
| connect=*net_service_name* | Specifies net service name of the database to which you want to connect. This is the net service name set in the `tnsnames.ora` file. |
| stop | Stops the OID Monitor process |

For example:

```
oidmon connect=dbsl stop
```

## Task 2: Start a Server Instance

Once the OID Monitor is running, start a server instance by using the OID Control Utility.

> **Note:** The value for the instance flag in the OID Control Utility should always be greater than or equal to one.

This section contains these topics:

- Starting an Oracle Directory Server Instance
- Stopping an Oracle Directory Server Instance
- Starting an Oracle Directory Replication Server Instance
- Stopping an Oracle Directory Replication Server Instance
- Restarting Directory Server Instances
- Troubleshooting Directory Server Instance Startup

## Starting an Oracle Directory Server Instance

The syntax for starting an Oracle directory server instance is:

```
oidctl connect=net_service_name server=oidldapd instance=server_instance_number
[configset=configset_number] [flags=' -p port_number -work maximum_number_of_
worker_threads_per_server -debug debug_level -l change_logging -server n'] start
```

| Argument | Description |
|---|---|
| `connect=net_service_name` | If you already have a `tnsnames.ora` file configured, this is the net service name specified in that file, located in `ORACLE_HOME/network/admin` |
| `server=oidldapd` | Type of server to start (valid values are OIDLDAPD and OIDREPLD). This is not case-sensitive. |
| `instance=server_instance_number` | Instance number of the server to start. Should be a number between 0 and 1000. |
| `configset=configset_number` | Configset number used to start the server. This defaults to `configset0` if not set. This should be a number between 0 and 1000. |
| `-p port_number` | Specifies a port number during server instance startup. Default port if not set is 389. |
| `-work maximum_number_of_worker_threads_per_server` | Specifies the maximum number of worker threads for this server |
| `-debug debug_level` | Specifies a debug level during Oracle directory server instance startup |
| `-l change_logging` | Turns replication change logging on and off. To turn it off, enter -l. To turn it on, omit the flag. The default is true (values = true and false). (directory server only) |
| `-server n` | Specifies the number of server processes to start on this port |
| `start` | Starts the server specified in the `server` argument. |

For example, to start an Oracle directory server instance whose net service name is dbs1, using configset5, at port 12000, with a debug level of 1024, an instance number 3, and in which change logging is turned off, type at the system prompt:

```
oidctl connect=dbs1 server=oidldapd instance=3 configset=5 flags='-p 12000
-debug 1024 -l ' start
```

When starting and stopping an Oracle directory server instance, the server name and instance number are mandatory. All other arguments are optional.

All keyword value pairs within the flags arguments must be separated by a single space.

Single quotes are mandatory around the flags.

The configset identifier defaults to zero (`configset0`) if not set.

---

**Note:** If you choose to use a port other than the default port (389 for non-secure usage or 636 for secure usage), you must tell the clients which port to use to locate the Oracle Internet Directory. If you use the default ports, clients can connect to the Oracle Internet Directory without referencing a port in their connect requests.

---

## Stopping an Oracle Directory Server Instance

OID Monitor must be running whenever you start or stop directory server instances.

At the system prompt, type:

```
oidctl connect=net_service_name server=OIDLDAPD instance=server_instance_number
stop
```

For example:

```
oidctl connect=dbs1 server=oidldapd instance=3 stop
```

## Starting an Oracle Directory Replication Server Instance

The syntax for starting the Oracle directory replication server is:

```
oidctl connect=net_service_name server=oidrepld instance=server_instance_number
[configset=configset_number] flags=' -h hostname -p port_number
-d debug_level -m [true | false]-z transaction_size ' start
```

| Argument | Description |
|---|---|
| `connect=net_service_name` | If you already have a `tnsnames.ora` file configured, then this is the name specified in that file, which is located in ORACLE_HOME/network/admin |
| `server=oidrepld` | Type of server to start (valid values are OIDLDAPD and OIDREPLD). This is not case-sensitive. |
| `instance=server_instance_number` | Instance number of the server to start. Should be a number between 0 and 1000. |
| `configset=configset_number` | Configset number used to start the server. This defaults to configset0 if not set. This should be a number between 0 and 1000. |
| `-p port_number` | Specifies a port number during server instance startup. Default port if not set is 389. |
| `-d debug_level` | Specifies a debug level during replication server instance startup |
| `-h` | Specifies the host name on which the server runs. (Replication server only) |
| `-m [true\|false]` | Turns conflict resolution on and off. The default is true (values = true and false). (Replication server only) |
| `-z transaction_size` | Specifies the number of changes applied in each replication update cycle. If you do not specify this, the number is determined by the Oracle directory server sizelimit parameter, which has a default setting of 1024. You can configure this latter setting. |
| `start` | Starts the server specified in the *server* argument. |

For example, to start the replication server with an instance=1, at port 12000, with debugging set to 1024, type at the system prompt:

```
oidctl connect=dbs1 server=oidrepld instance=1 flags='-p 12000 -h eastsun11 -d
1024' start
```

When starting and stopping an Oracle directory replication server, the `-h` flag, which specifies the host name, is mandatory. All other flags are optional.

All keyword value pairs within the flags arguments must be separated by a single space.

Single quotes are mandatory around the flags.

The configset identifier defaults to zero (`configset0`) if not set.

> **Note:** If you choose to use a port other than the default port (389 for non-secure usage or 636 for secure usage), you must tell the clients which port to use to locate the Oracle Internet Directory. If you use the default ports, clients can connect to the Oracle Internet Directory without referencing a port in their connect requests.

## Stopping an Oracle Directory Replication Server Instance

OID Monitor must be running whenever you start or stop directory server instances.

At the system prompt, type:

```
oidctl connect=net_service_name server=OIDREPLD instance=server_instance_number
stop
```

For example:

```
oidctl connect=dbs1 server=oidrepld instance=1 stop
```

## Restarting Directory Server Instances

To restart a directory server instance, at the system prompt, type:

```
oidctl connect=net_service_name server={oidldapd|oidrepld} instance=server_
instance_number  restart
```

OID Monitor must be running whenever you start, stop, or restart directory server instances.

If you try to contact a server that is down, you receive from the SDK the error message `81—LDAP_SERVER_DOWN`.

If you change a configuration set entry that is referenced by an active server instance, you must stop that instance and restart it to effect the changed value in the configuration set entry on that server instance. You can either issue the STOP command followed by the START command, or you can use the RESTART command. RESTART both stops and restarts the server instance.

For example, suppose that Oracle directory server instance1 is started, using configset3, and with the net service name dbs1. Further, suppose that, while instance1 is running, you change one of the attributes in configset3. To enable the change in configset3 to take effect on instance1, you enter the following command:

```
oidctl connect=dbs1 server=oidldapd instance=1 restart
```

If there are more than one instance of the Oracle directory server running on that node using configset3, then you can restart all the instances at once by using the following command syntax:

```
oidctl connect=dbs1 server=oidldapd restart
```

Note that this command restarts all the instances running on the node, whether they are using configset3 or not.

---

**Important Note:** During the restart process, clients cannot access the Oracle directory server instance. However, the process takes only a few seconds to execute.

---

## Troubleshooting Directory Server Instance Startup

If the directory server fails to start, you can override all user-specified configuration parameters to start the directory server and then return the configuration sets to a workable state by using the ldapmodify operation.

To start the directory server using its hard-coded default parameters instead of the configuration parameters stored in the directory, type at the system prompt:

```
oidctl connect=net_service_name flags='-p port_number -f'
```

The `-f` option in the flags starts the server with hard-coded configuration values, overriding any defined configuration sets except for the values in `configset0`.

## Task 3: Reset the Default Security Configuration

When you first install Oracle Internet Directory, the default configuration grants to all users read, browse, and search access to all entries in the directory. One of the first things you need to do is establish and implement an access control policy to ensure that each user receives the appropriate authorization. Oracle Corporation specifically recommends that you control access to the subentry `subSchemaSubEntry` and its children because these objects contain information about the directory.

Moreover, when you load directory entries, you are creating a hierarchy of directory entries. You must therefore establish:

- Permissions to load entries into this hierarchy
- Directory access for clients that need read, modify, and write access to the directory entries

**See Also:**

- Chapter 9, "Managing Directory Access Control" for a detailed explanation of access control options and instructions for setting up security
- Chapter 4, "Using the Administration Tools" for information about the administration tools you use to configure security
- Appendix E, "Schema Elements" for syntax and usage notes for the command line tools

# Upgrading from an Earlier Release of Oracle Internet Directory

Oracle Internet Directory release 2.1.1 allows you to upgrade from either Oracle Internet Directory release 2.0.4.x or release 2.0.6. You choose to upgrade to a release 2.1.1 when prompted during the installation process.

In a replicated environment, a node running release 2.1.1 can co-exist with nodes running previous releases of Oracle Internet Directory. Moreover, in a replicated environment, upgrade of one node to release 2.1.1 requires no network downtime. The other nodes can remain available while the upgrade progresses.

This section contains these topics:

- Upgrading in a Single Node Environment
- Upgrading in a Multi-Node Environment

## Upgrading in a Single Node Environment

To upgrade on a single node, follow the instructions in the installation documentation for your operating system.

## Upgrading in a Multi-Node Environment

Upgrading a multi-node Oracle Internet Directory system to release 2.1.1 requires special attention. This section discusses the two ways to upgrade a multi-node Oracle Internet Directory system:

- Upgrading one node at a time
- Upgrading all the nodes at the same time

### Upgrading One Node at a Time

Use this method if you do not want any system downtime. While the upgrade on one node is in progress, it allows all the other nodes to remain available. However, using this method requires that you clearly understand and strictly follow these guidelines:

- When you are upgrading a replication network one node at a time, the upgrade is not complete until all nodes are upgraded. However, during this period, all network nodes, except the one being upgraded, remain available. This is considered a transient state and, to indicate this, the attribute `orclupgradeinprogress` in the DSE root is set to `TRUE`. This attribute is created during the upgrade procedure.
- When all the nodes in a replicated Oracle Internet Directory network have been upgraded to release 2.1.1, the system is no longer in a transient state. At this point, set the attribute `orclupgradeinprogress` to `FALSE` on all the nodes.
- During the transient state, that is, when the `orclupgradeinprogress` attribute in the DSE root is set to `TRUE`, the release 2.1.1 node performs special processing on change log entries it generates for other nodes. This is required for backward compatibility.
- During the transient state, do not use the new password encryption scheme. Doing this creates inconsistencies in password values across various nodes, and nodes running previous releases will disable the authentication. Instead of doing this, continue using the existing password encryption scheme. Once the entire network is upgraded, you can start using the new password encryption scheme.

- While the upgrade is going on, only one node should be Read-Write. The rest should be Read-Only.
- Do not perform binary attribute modification on the upgraded node. Such modifications fail on 2.0.4.x and 2.0.6 nodes.
- Be sure to perform the upgrade on the **master definition site (MDS)** before you upgrade the master sites.

**See Also:**

- "What's New in Oracle Internet Directory?" on page xxix for a list of encryption algorithms supported for passwords in Oracle Internet Directory release 2.1.1
- "Post Upgrade Procedure for Password Encryption" on page 3-18

Perform the following tasks, first on the MDS, then on the master sites.

#### Task 1: Stop the Oracle Directory Replication Server on the Node to be Upgraded

**See Also:** "Stopping an Oracle Directory Replication Server Instance" on page 3-7

#### Task 2: Stop the Oracle Directory Server on the Node to be Upgraded

**See Also:** "Stopping an Oracle Directory Server Instance" on page 3-5

#### Task 3: Stop OID Monitor on the Node to be Upgraded

**See Also:** "Stopping the OID Monitor" on page 3-3

**Task 4: Delete Jobs on Other Nodes** Before shutting down the database at the MDS, run the script delasrjobs.sql located in `/oidupgrade/`on the installation CD. This script deletes **ASR** jobs on other master sites that push changes to the MDS. Deleting these jobs temporarily removes the MDS from the replication environment so that no changes can be applied to it. Other nodes, however, remain operational and continue replicating changes.

**Task 5: Shutdown Database and Listener on the Node to be Upgraded** If you do not shutdown the database and listener, then Oracle Universal Installer prompts you to do it.

**See Also:**

- *Net8 Administrator's Guide* for instructions on stopping the listener
- *Oracle8i Administrator's Guide* for instructions on shutting down the database server

**Task 6: Upgrade the Node to Oracle Internet Directory Release 2.1.1** Run Oracle Universal Installer to upgrade to Oracle Internet Directory release 2.1.1, which uses Oracle8*i* release 8.1.7. The installer both migrates the database and upgrades Oracle Internet Directory.

**Task 7: Start the Database and Listener** After the upgrade is completed, make sure that the database and listener are up and running.

Test the connectivity to other nodes. If connectivity is broken, then use the backup copies of `listener.ora`, `sqlnet.ora` and `tnsnames.ora` and restart the listener. The backup files are named `listener`*`date`*`.bak`, `sqlnet`*`date`*`.bak` and `tnsnames`*`date`*`.bak`.

**Task 8: Create Push Jobs on Other Nodes** After you have upgraded the node, create jobs on other nodes. You do this by executing `$`*`ORACLE_HOME`*`/ldap/admin/creasrjobs.sql` on the upgraded node. This script creates on the other nodes the jobs that were deleted in "Task 4: Delete Jobs on Other Nodes" on page 3-11. These jobs now start pushing the existing changes and new changes on other nodes to the node you have just upgraded.

**Task 9: Perform Post Upgrade Procedure for Password Encryption** Once the node is upgraded, perform the post upgrade procedure for password encryption as described in "Post Upgrade Procedure for Password Encryption" on page 3-18.

**Task 10: Start OID Monitor**

**See Also:** "Starting the OID Monitor" on page 3-2

**Task 11: Start the Oracle Directory Server**

> **See Also:** "Starting an Oracle Directory Server Instance" on page 3-4

**Task 12: Start the Oracle Directory Replication Server**

> **See Also:** "Starting an Oracle Directory Replication Server Instance" on page 3-6

**Task 13: Upgrade Other Master Sites** After upgrading the MDS, upgrade other master sites one at a time. Perform tasks 1 through 12 on each master site until all the nodes are upgraded.

**Task 14: Update the orclupgradeinprogress Attribute on All the Nodes** After all the nodes are upgraded to Oracle Internet Directory release 2.1.1, modify the `orclupgradeinprogress` attribute to `FALSE` on all the nodes. To do this:

1. Edit the input file as follows:

```
dn:
modify:replace
replace:orclupgradeinprogress
orclupgradeinprogress:FALSE
```

2. Use ldapmodify to load the file:

```
ldapmodify -D "cn=orcladmin" -w welcome -h host_name -p  port_number -f
input_file.ldif
```

> **See Also:** Chapter 10, "Managing Directory Replication" for information about the MDS

### Upgrading All the Nodes at the Same Time

Use this method to upgrade all the nodes at the same time. If you use this method, then the system is unavailable during the upgrade process.

#### Task 1: Set All the Nodes in the Network to Read-Only Mode

1. Edit the input file as follows:

```
dn:
changetype:modify
replace:orclservermode
orclservermode:r
```

2. Run the following command against all the nodes in the replication network:

```
ldapmodify -D "cn=orcladmin" -w welcome -h host_name -p port_number -f
input_file.ldif
```

**Task 2: Wait Until All the Changes in the Change Log Queue Have Been Applied** Before moving to next step, wait for the change log queue to empty. If you skip this step, then changes in the change log queue will be applied once nodes are upgraded.

#### Task 3: Stop the Oracle Directory Replication Server on All Nodes

> **See Also:** "Stopping an Oracle Directory Replication Server Instance" on page 3-7

#### Task 4: Stop the Oracle Directory Server on All Nodes

> **See Also:** "Stopping an Oracle Directory Server Instance" on page 3-5

#### Task 5: Stop OID Monitor on All Nodes

> **See Also:** "Stopping the OID Monitor" on page 3-3

**Task 6: Shutdown the Database and the Listener on All Nodes** If you do not shutdown the database and listener, then Oracle Universal Installer prompts you to do it.

**See Also:**

- *Net8 Administrator's Guide* for instructions on stopping the listener
- *Oracle8i Administrator's Guide* for instructions on shutting down the database server

**Task 7: Upgrade All the Nodes to Oracle Internet Directory Release 2.1.1** Run Oracle Universal Installer to upgrade to Oracle Internet Directory release 2.1.1, which uses Oracle8*i* release 8.1.7. The installer both migrates the database and upgrades Oracle Internet Directory.

**Task 8: Start the Database and Listener on All Nodes** After the upgrade is completed, make sure that the database and listener are up and running.

Test the connectivity to other nodes. If connectivity is broken, then use the backup copies of `listener.ora`, `sqlnet.ora` and `tnsnames.ora` and restart the listener. The backup files are named `listener`*`date`*`.bak`, `sqlnet`*`date`*`.bak` and `tnsnames`*`date`*`.bak`.

**Task 9: Perform Post Upgrade Procedure for Password Encryption** Once the node is upgraded, perform the post upgrade procedure for password encryption as described in "Post Upgrade Procedure for Password Encryption" on page 3-18.

**Task 10: Start OID Monitor on All Nodes**

**See Also:** "Starting the OID Monitor" on page 3-2

**Task 11: Start Oracle Directory Server on All Nodes**

**See Also:** "Starting an Oracle Directory Server Instance" on page 3-4

**Task 12: Start Oracle Directory Replication Server on All Nodes**

**See Also:** "Starting an Oracle Directory Replication Server Instance" on page 3-6

**Task 13: Update the orclupgradeinprogress Attribute on All the Nodes** After all the nodes are upgraded to Oracle Internet Directory release 2.1.1, modify the `orclupgradeinprogress` attribute to `FALSE` on all the nodes. To do this:

1. Edit the input file as follows:

```
dn:
modify:replace
replace:orclupgradeinprogress
orclupgradeinprogress:FALSE
```

2. Use ldapmodify to load the file:

```
ldapmodify -D "cn=orcladmin" -w welcome -h host_name -p  port_number -f
input_file.ldif
```

Perform this modification on all the nodes in the replication environment

## LDIF-Based Upgrading

Oracle Corporation recommends that you use the LDIF-based backup procedure to backup your existing release Oracle Internet Directory. This is explained in this section.

Normally, you do not need to perform LDIF-based upgrading. Use this method when you cannot successfully run the database-based upgrade process.

The LDIF-based upgrade process requires the following procedures on a node being upgraded:

**Task 1: Backup the Older Version of Oracle Internet Directory** Be sure that the Oracle directory server is not running, then run the script backup_oid.sh located in the `/oidupgrade` directory on the CD.

The syntax to run backup_oid.sh is:

```
backup_oid.sh -connect net_service_name -pass password_for_DB_account_'ods'
```

The backup_oid.sh script does the following:

- Exports Oracle Internet Directory schema. As it does this, it generates `.dmp` files—for example, `attr_store.dmp`—in $*ORACLE_HOME*`/ldap/load` directory
- Backs up the OID subtree using the Ldifwrite utility. As it does this, it generates the file `OID_userdata.ldif` in $*ORACLE_HOME*`/ldap/load`. The subtree

under `cn=OracleSchemaVersion` (if it exists) is also backed up as `orcl_schemaver.ldif` in the $*ORACLE_HOME*`/ldap/load` directory.

If you plan to install Oracle Internet Directory release 2.1.1 in the same *ORACLE_HOME*, then save these generated files in some other location.

**Task 2: Perform a Fresh Installation of Oracle Internet Directory Release 2.1.1**

> **See Also:** Installation documentation for your operating system

**Task 3: Restore the User-Defined Schema and Data from the Previous Version of Oracle Internet Directory:** To do this:

1. Make sure that the Oracle directory server is not running.
2. Copy the following files to $*ORACLE_HOME*`/ldap/load`:
   - Backed up Oracle Internet Directory schema dump files—that is, files with the extension `.dmp`
   - The file `OID_userdata.ldif`
3. Run the script restore_oid.sh located in $*ORACLE_HOME*`/ldap/install`.

   The syntax for restore_oid.sh is:

   `restore_oid.sh -connect` *net_service_name* `-pass` *password_for_DB_account_'ods'*

The restore_oid.sh script does the following:

- Imports the Oracle Internet Directory schema from the dump files
- Inserts the schema differences between the previous release—that is, either 2.0.6 or 2.0.4—and 2.1.1
- Bulkloads the data from the LDIF file with the `-restore` option

**Task 4: Upgrade Passwords** Run the cryptupgrd.sh script located in $*ORACLE_HOME*`/ldap/bin/` to upgrade passwords.

The syntax for cryptupgrd.sh is:

`cryptupgrd.sh -connect` *net_service_name* `-pass` *password_for_DB_account_'ods'*

### Post Upgrade Procedure for Password Encryption

In release 2.0.6 and release 2.0.4, the user password was encrypted by using only one encryption algorithm, namely MD4. A flag in the root DSE, namely, `orcluseencrypt`, merely toggled encryption on and off. By contrast, Oracle Internet Directory release 2.1.1 supports multiple hash schemes.

Oracle Internet Directory release 2.1.1 stores the hash scheme as a prefix to the password value. A new attribute created in the root DSE during the upgrade procedure indicates the default hashing scheme. If the supplied password is not already encrypted, then the Oracle directory server uses this default value to encrypt passwords.

The post upgrade procedure for password encryption adds a prefix `MD4` to all the existing password values in the directory. The time taken by this procedure to finish varies depending on the number of entries in the directory.

To run this procedure, enter the following command:

```
cryptupgrd.sh -connect  net_service_name -pass password_for_DB_account_'ods'
```

**See Also:**

- "What's New in Oracle Internet Directory?" on page xxix for a list of encryption algorithms supported for passwords in Oracle Internet Directory release 2.1.1

# 4

# Using the Administration Tools

This chapter introduces the various administration tools of Oracle Internet Directory. It discusses the online administration tool, called Oracle Directory Manager, and tells you how to launch it, navigate through it, and connect to directory servers with it. It also introduces the command line and bulk tools.

This chapter contains these topics:

# Using Oracle Directory Manager

Oracle Directory Manager is a Java-based tool for administering Oracle Internet Directory. This section describes some of its basic features. More specific instructions are found in sections throughout this book that explain how to perform various tasks.

This section contains these topics:

- Starting Oracle Directory Manager
- Connecting to a Directory Server
- Navigating Oracle Directory Manager
- Connecting to Additional Directory Servers
- Disconnecting from a Directory Server
- Performing Administration Tasks by Using Oracle Directory Manager

## Starting Oracle Directory Manager

Before you can launch Oracle Directory Manager, you must have a directory **server instance** running.

**See Also:**

- Chapter 3, "Preliminary Tasks" for instructions on running a server instance
- "Oracle Internet Directory Architecture" on page 2-20 for a conceptual explanation of directory server instances

To start Oracle Directory Manager, follow the instructions for your operating system:

| Operating System | Instructions |
|---|---|
| Windows NT or Windows 95 | From the Start menu, click Programs > *ORACLE_HOME* > Oracle Internet Directory > Oracle Directory Manager |
| Sun Solaris | If you have not set the path, then navigate to `ORACLE_HOME/bin`.<br>Type at the system prompt:<br>`oidadmin` |

The first time you start Oracle Directory Manager, an alert tells you that you must connect to a server. Click OK.

## Connecting to a Directory Server

To connect to a directory server:

1. In the Directory Server Connection dialog box, type the name and port number of an available server.

   The default port is 389. You can change the port if you wish. However, if you have an Oracle directory server running on a non-default port, be sure that any clients that use that server are informed of the correct port.

   Click OK. The Oracle Directory Manager Connect dialog box appears.

2. In each field of the Credentials tab page, type the information specific to this server instance as described in the next table.

| Field | Description |
|---|---|
| User | The first time you log in, do so either as the **superuser** or anonymously. If you intend to configure SSL features during this session, login as the super user.<br>If you are logging in as the super user, in the User box, type `cn=orcladmin`.<br>If you are logging in anonymously, leave the User box empty.<br>If you have already set up the user's entry by using LDAP command line tools, you can enter that user's entry in one of two ways:<br>■ Browse and select that entry by using the button to the right of the User field<br>■ Type the **distinguished name (DN)** for that user's entry by using the correct format, for example,<br>`cn=Susie Brown,ou=HR,o=acme,c=us` |

| Field | Description |
|---|---|
| Password | If you are logging in as the super user and you specified a password for the super user during installation, in the Password box, type the password you specified. Otherwise, type the default password, namely, `welcome`. After you are logged into Oracle Directory Manager and have connected to a directory server, you should change this password to protect the directory.<br><br>If you are logging in anonymously, leave the Password box empty.<br><br>If you want to login as a specific directory user, enter the corresponding password.<br><br>**See Also:** "Managing Super, Guest, and Proxy Users" on page 5-20 for instructions on how to change the password |
| Server | From the Server list, select the host containing the directory server to which you want to connect.<br><br>If you are already connected to a directory server, and you want to connect to a directory server on a different host:<br><br>**1.** Click the button to the right of the Server field. A dialog box displays a list of available servers.<br><br>**2.** Select a server.<br><br>**3.** Click OK.<br><br>To add a directory server:<br><br>**1.** Click Add. The Directory Server Connection dialog box appears.<br><br>**2.** Type the name of the directory server you want to add.<br><br>**3.** Click OK. |
| Port | The default port (389) appears in this field. If there is more than one directory server instance on the same host, each directory server instance has a different port, and that port number appears in this field when you select the directory server instance.<br><br>To change this port number:<br><br>**1.** Click the button to the right of the Server field.<br><br>**2.** In the Select Directory Server dialog box, select the directory server.<br><br>**3.** Click Edit. The Directory Server Connection dialog box appears.<br><br>**4.** In the Directory Server Connection dialog box, in the Port field, enter the new port number, then click Ok. |

| Field | Description |
|---|---|
| SSL Enabled | Selecting this check box causes all commands you issue by using Oracle Directory Manager to be sent over Secure Sockets Layer (SSL).<br><br>You can connect to a directory server either with or without SSL. If you connect by using SSL, then Oracle Directory Manager becomes an SSL client.<br><br>You can connect in this way if both of the following two conditions are met:<br><br>■ The server to which you are connecting uses SSL. If that server does not use SSL, and you select this check box, then authentication will fail when you try to connect.<br><br>■ You have already created a wallet containing a certificate and a list of trusted certificates. |

**See Also:**

- Chapter 8, "Managing Secure Sockets Layer (SSL)" for instructions on enabling SSL
- Appendix C for instructions on creating a wallet
- "Entries" on page 2-2 for instructions on formatting distinguished names
- "Configuring SSL Parameters" on page 8-2 for information about changing ports and their impact on security

3. If you selected the SSL Enabled check box on the Credentials tab, then select the SSL tab.

**4.** Enter the requested data in the fields as described in the next table.

| Field | Description |
|---|---|
| SSL Location | If the user's wallet is on the local machine, then type the wallet path and file name by using this syntax:<br>`file: absolute_path_name`<br>If the wallet is on another machine, link to that location, then enter the linked path and file name of the wallet. |
| SSL Password | The password to open the user's wallet |
| SSL Authentication | Select the authentication level:<br>■ No SSL Authentication—Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. If you selected the SSL Enabled check box on the Credentials tab, and choose this option, then only SSL encryption / decryption will be used.<br>■ SSL Client and Server Authentication—Two-way authentication. Both client and server send certificates to each other.<br>■ SSL Server Authentication—One-way authentication. Only the directory server authenticates itself to the client by sending its certificate to the client. |

**Note:** If the server requires two-way authentication, then each Oracle Directory Manager user must have a unique wallet. If one-way authentication is specified, then several Oracle Directory Manager users can use a single wallet.

**5.** Click Login. Oracle Directory Manager appears.

## Navigating Oracle Directory Manager

This section provides an overview of Oracle Directory Manager, and explains the items in the menu bar and the buttons on the toolbar.

### Overview of Oracle Directory Manager

Like the directory itself, the navigator pane (left side of the double window interface) has a tree-like structure. When Oracle Directory Manager first opens, the navigator pane shows only one tree item, Oracle Internet Directory Servers. By clicking the plus sign(+) next to the tree item, subcomponents of that tree item appear.

In the right pane, some windows contain buttons labeled Apply and OK. If you press Apply, the changes you have made are committed, and the window remains available for more changes. If you press OK, the changes you have made are committed, and the window closes.

Similarly, some windows have buttons that are labeled Revert and Cancel. If you press Revert, the changes you have made in that window do not take effect, and the window stays open for further work. If you press Cancel, the changes you have made in that window do not take effect, and the window closes.

### The Oracle Directory Manager Menu Bar

The next table lists and describes the menus you can access by using the menu bar. Menu items become enabled or disabled depending on the pane or tab page you are displaying.

| Menu | Menu Items |
|---|---|
| File | Create—Adds an object |
| | Create Like—Adds a new object by using the object selected in the navigator pane as a template |
| | Connect—Connects to a directory server selected in the navigator pane |
| | Disconnect—Disconnects from a directory server selected in the navigator pane |
| | Exit—Exits Oracle Directory Manager |
| Edit | Edit—Modifies an object |
| | Remove—Removes a selected object |
| | Find Object Classes—Searches for an object class |

| Menu | Menu Items |
|---|---|
| View | Refresh—Updates data stored in memory to reflect changes in the database<br><br>Tear-Off—Generates a secondary dialog containing the fields and values displayed in Oracle Directory Manager's right pane. This is useful when comparing two pieces of information. |
| Operations | Create Object Class—Displays the New Object Class dialog box that you use to add a new object class<br><br>Create Attribute—Displays the New Attribute Type dialog box that you use to add a new attribute to an entry<br><br>Create Access Ctrl Point—Displays the New Access Control Point dialog box that you use to add a new **Access Control Policy Point**.<br><br>Create Entry—Displays the New Entry dialog box that you use to add a new directory entry<br><br>Refresh Entry—Updates data for entries stored in memory to reflect changes in the database<br><br>Refresh Subtree Entries—Updates the children of entries stored in memory to reflect changes in the database<br><br>Drop Index—Removes an index from an attribute. When you select this item, an alert asks you to confirm that you want to drop the index.<br><br>Search ACPs—Enables you to configure ACP searches<br><br>User Preferences—Displays a dialog box that enables you to:<br><br>■ Configure the display of entry search results<br><br>■ Establish whether ACPs are displayed whenever Oracle Directory Manager runs, or only as the result of a search |
| Help | Contents—Displays the Contents tab page of the Help navigator<br><br>Search for Help On...—Displays the Help Search dialog box that you use to search for words in the online help guide<br><br>About Oracle Internet Directory—Displays Oracle Internet Directory version information |

## The Oracle Directory Manager Toolbar

Figure 4–1 and the accompanying table illustrate and describe the Oracle Internet Directory toolbar. Buttons become enabled or disabled depending on the pane or tab page you are displaying in Oracle Directory Manager.

***Figure 4–1 Oracle Directory Manager Toolbar***

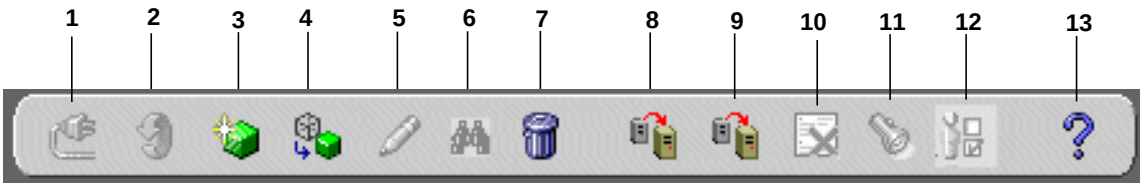


| Button | Purpose |
|---|---|
| 1 | Connect / Disconnect—Connects to or disconnect from a directory server selected in the navigator pane |
| 2 | Refresh—Updates data for objects other than entries that are stored in memory to reflect changes in the database |
| 3 | Create—Adds a new object |
| 4 | Create Like—Adds a new object by using another object as a template |
| 5 | Edit—Modifies an object |
| 6 | Find Object Classes—Searches for an object class |
| 7 | Delete—Removes an object |
| 8 | Refresh Entry—Updates data for entries stored in memory to reflect changes in the database |
| 9 | Refresh SubTree Entries—Updates the children of entries stored in memory to reflect changes in the database |
| 10 | Drop Index—Removes an index from an attribute. When you click this button, an alert asks you to confirm that you want to drop the index. |
| 11 | Search—Enables you to configure ACP searches |
| 12 | User Preferences—Enables you to configure the display of ACPs in the navigator pane, as well as entries in a search operation |
| 13 | Help—Displays the Help system |

## Connecting to Additional Directory Servers

You can connect to more than one directory server at a time, and then view and modify the data, schema, and security for each directory server. If you do this, then each server is listed in the navigator pane under Oracle Internet Directory Servers.

To connect to an additional directory server:

1. In the navigator pane, select Oracle Internet Directory Servers.
2. In the right pane, click New.
3. Follow the login procedures described in "Connecting to a Directory Server" on page 4-3.

## Disconnecting from a Directory Server

To disconnect from a directory server by using Oracle Directory Manager, choose File > Disconnect. Also, when you exit Oracle Directory Manager, connections between all directory servers and the directory are automatically disconnected.

All connection information is stored in the user's home directory in the file `osdadmin.ini`.

When you restart Oracle Directory Manager, all previously connected server connections appear in the Directory Server Login dialog box.

## Performing Administration Tasks by Using Oracle Directory Manager

You can perform most of the Oracle Internet Directory administrative tasks through Oracle Directory Manager. Tasks that you cannot perform through Oracle Directory Manager involve running processes, such as starting and stopping the OID Monitor (oidmon) process and starting and stopping server instances. To perform tasks that you cannot perform with Oracle Directory Manager, use the appropriate LDAP command line tool.

The following table lists the task areas managed by Oracle Directory Manager and where to find instructions for using it in each area.

| Task Area | Instructions |
|---|---|
| Schema administration | "Managing Object Classes by Using Oracle Directory Manager" on page 6-6<br>"Managing Attributes by Using Oracle Directory Manager" on page 6-17 |
| Entries management | "Managing Entries by Using Oracle Directory Manager" on page 7-2 |
| ACP administration | "Managing Access Control by Using Oracle Directory Manager" on page 9-16 |
| Partitioning and replication | Chapter 10, "Managing Directory Replication" |

## Using Command Line Tools

Oracle Internet Directory provides several command line tools for manipulating directory entries and attributes. This section explains the kind of tasks you can perform with each tool.

The command line tools act on objects that are in text files written in the LDAP Data Interchange Format (LDIF).

> **See Also:** "LDAP Data Interchange Format (LDIF) Syntax" on page A-2 for information on formatting an LDIF file

The following table lists each command line tool, the task(s) you can perform with it, and where to find syntax and usage notes.

| Tool | Task(s) | Syntax and Usage Notes |
|---|---|---|
| ldapsearch | Search for directory entries. | "ldapsearch Syntax" on page A-18 |
| ldapbind | Authenticate user / client to a directory server. | "ldapbind Syntax" on page A-8 |
| ldapadd | Add entries one at a time. | "ldapadd Syntax" on page A-4 |
| ldapaddmt | Add several entries concurrently by using this multithreaded tool. | "ldapaddmt Syntax" on page A-6 |
| ldapmodify | Create, update, and delete attribute data for an entry. | "ldapmodify Syntax" on page A-13 |
| ldapmodifymt | Modify several entries concurrently by using this multithreaded tool. | "ldapmodifymt Syntax" on page A-16 |
| ldapdelete | Delete entries. | "ldapdelete Syntax" on page A-10 |
| ldapcompare | See whether an entry contains a specified attribute value. | "ldapcompare Syntax" on page A-9 |
| ldapmoddn | Modify the DN or RDN of an entry, rename an entry or a subtree, or move an entry or a subtree under a new parent. | "ldapmoddn Syntax" on page A-11 |

**See Also:** "Using NLS with Command Line Tools" on page 12-5 for a discussion of command line tools and NLS

# Using Bulk Tools

Bulk tools enable you to create and manage large numbers of directory entries from data residing in, or created by, other applications.

---

**Important Note:** To use these tools you must provide the Oracle Internet Directory password. The default password is `ods`, although the system administrator can change it by using the OID Database Password Utility.

---

**See Also:**

- "Using the OID Database Password Utility" on page 4-14
- "OID Database Password Utility Syntax" on page A-37

The table that follows lists each bulk tool, the task(s) you can perform with it, and where to find syntax and usage notes.

| Tool | Task(s) | Syntax and Usage Notes |
|---|---|---|
| bulkload | Load large number of entries to Oracle Internet Directory through LDIF files | "bulkload Syntax" on page A-23 |
| ldifwrite | Copy data from the directory information base into an LDIF file that can be read by any LDAP compliant directory server. You can use ldifwrite in conjunction with bulkload. You can also use ldifwrite to back up information from all or part of a directory. | "ldifwrite Syntax" on page A-27 |
| bulkmodify | Modify a large number of existing entries efficiently | "bulkmodify Syntax" on page A-25 |
| bulkdelete | Delete a subtree efficiently | "bulkdelete Syntax" on page A-22 |

## Using OID Control Utility

OID Control Utility is a command line tool for starting and stopping the server. The commands are interpreted and executed by the OID Monitor process.

**See Also:**

- "OID Control Utility Syntax" on page A-31
- "Oracle Internet Directory Architecture" on page 2-20 for a conceptual description

## Using the Catalog Management Tool

Oracle Internet Directory uses indexes to make attributes available for searches. When Oracle Internet Directory is installed, the entry `cn=catalogs` lists available attributes that can be used in a search. Only those attributes that have an equality matching rule can be indexed.

If you want to use additional attributes in search filters, you must add them to the catalog entry. You can do this at the time you create the attribute by using Oracle Directory Manager. However, if the attribute already exists, then you can index it only by using the Catalog Management tool.

**See Also:**

- "Catalog Management Tool Syntax" on page A-28 for syntax and usage notes
- "Indexing an Attribute by Using Command Line Tools" on page 6-29
- "Indexing an Attribute When You Create It" on page 6-27

## Using the OID Database Password Utility

Oracle Internet Directory uses a password when connecting to an Oracle database. The default for this password when you install Oracle Internet Directory is `ODS`. You can change this password by using the OID Database Password Utility.

**See Also:** "OID Database Password Utility Syntax" on page A-37 for syntax and usage notes

## Using the Replication Tools

When a replication conflict arises, Oracle directory replication server places the change in the retry queue and tries to apply it from there for a specified number of times. If it fails after that specified number, then the replication server puts the change in the human intervention queue. From there, the replication server repeats the change application process at less frequent intervals while awaiting your action.

At this point, you need to:

1. Examine the change in the human intervention queues
2. Reconcile the conflicting changes
3. Place the change either back into the retry queue or into the purge queue.

The following replication tools assist you in this process:

| | |
|---|---|
| OID reconciliation tool | Enables you to synchronize conflicting changes |
| Human intervention queue manipulation tool | Enables you to move changes from the human intervention queue to either the retry queue or the purge queue |

**See Also:**

- "Using the OID Reconciliation Tool" on page 10-33
- "Using the Human Intervention Queue Manipulation Tool" on page 10-30

## Using the OID Database Statistics Collection Tool

The OID database statistics collection tool, located in `$ORACLE_HOME/ldap/admin/`, assists in capacity planning. It helps you analyze the various database `ods` schema objects so that you can estimate the statistics.

**See Also:** "OID Database Statistics Collection Tool Syntax" on page A-37

# Administration Tasks at a Glance

Oracle Internet Directory administration tasks are described throughout this manual. Table 4–1 points you to the information you need for some of the more common tasks.

***Table 4–1 Common Administration Tasks and Where To Find Instructions***

| Task | Information |
|---|---|
| **Managing Attributes** | |
| Add, modify, or delete an attribute by using command line tools | "Managing Attributes by Using Command Line Tools" on page 6-28 |
| Add, modify, or delete an attribute by using the Oracle Directory Manager | "Managing Attributes by Using Oracle Directory Manager" on page 6-17 |
| **Managing Entries** | |
| Add, modify, or delete a directory entry by using command line tools | "Managing Entries by Using Command Line Tools" on page 7-11 |
| Add, modify, or delete a directory entry by using Oracle Directory Manager | "Managing Entries by Using Oracle Directory Manager" on page 7-2 |
| Import bulk data files | "bulkload Syntax" on page A-23<br>"LDAP Data Interchange Format (LDIF) Syntax" on page A-2 |
| View Directory Information Tree (DIT) hierarchy of entries | "Managing Entries by Using Oracle Directory Manager" on page 7-2 |
| **Managing Object Classes** | |
| Add, modify, or delete object classes by using command line tools | "Managing Object Classes by Using Command Line Tools" on page 6-14 |
| Add, modify, or delete object classes by using Oracle Directory Manager | "Managing Object Classes by Using Oracle Directory Manager" on page 6-6 |
| **Managing Security** | |
| Set up an Access Control Policy Point (ACP) | Chapter 9, "Managing Directory Access Control" |
| Set up security | Chapter 8, "Managing Secure Sockets Layer (SSL)" |
| **Managing Servers** | |
| Configure server instance parameters by using command line tools | "Managing Server Configuration Set Entries by Using Command Line Tools" on page 5-10 |

*Table 4–1 Common Administration Tasks and Where To Find Instructions*

| Task | Information |
|---|---|
| Configure server instance parameters by using the Oracle Directory Manager | "Managing Server Configuration Set Entries by Using Oracle Directory Manager" on page 5-4 |
| Connect to a directory by using Oracle Directory Manager | "Connecting to a Directory Server" on page 4-3<br>"Connecting to Additional Directory Servers" on page 4-10 |
| Run the directory server processes | Chapter 3, "Preliminary Tasks" |
| Stop the directory server processes | Chapter 3, "Preliminary Tasks" |
| View system operational attributes | "Setting System Operational Attributes by Using Oracle Directory Manager" on page 5-14<br>"Setting System Operational Attributes by Using ldapmodify" on page 5-15 |
| **Managing Replication** | |
| Set up replication | Chapter 10, "Managing Directory Replication" |
| Resolve replication change conflicts | "Conflict Resolution in Replication" on page 2-31 |
| Move replication changes from human intervention queue to either the retry queue or the purge queue | "Using the Human Intervention Queue Manipulation Tool" on page 10-30 |

# Part II

## Managing Oracle Internet Directory

Part II guides you through the tasks required to configure and maintain Oracle Internet Directory. This part contains these chapters:

# 5

# Managing an Oracle Directory Server

This chapter explains how to manage an Oracle directory server by using Oracle Directory Manager and command line tools.

This chapter contains these topics:

- Managing Server Configuration Set Entries
- Setting System Operational Attributes
- Managing Naming Contexts
- Managing Password Encryption
- Configuring Searches
- Managing Super, Guest, and Proxy Users
- Setting Debug Logging Levels
- Using Audit Log
- Viewing Active Server Instance Information
- Changing the Password to an Oracle Data Server

**See Also:** Chapter 3, "Preliminary Tasks" for instructions on starting and stopping directory server instances

# Managing Server Configuration Set Entries

When you start an Oracle directory server by using the **OID Control Utility**, that start message refers to a **configuration set entry** containing server parameters. You can add, modify, and delete configuration set entries by using either Oracle Directory Manager or the appropriate command line tool.

**See Also:**

- "Configuration Set Entries" on page 2-24 for a conceptual overview of configuration set entries
- "Task 2: Start a Server Instance" on page 3-3 for instructions on how to start the server by using OID Control Utility

This section contains these topics:



## Preliminary Considerations

Although you can change values in the default configuration set, namely, `configset0`, all of your changes will be carried over to every new configuration set entry that you create. This is because `configset0` values are used as the template for all new configuration set entries.

When you want to change values that should not always be in effect for every instance of the server that you run, it is better to create new configuration set entries. Note that, in release 2.1.1, this applies to the Oracle directory server instances only. The Oracle replication directory server supports only one configuration set in this release.

You may want to establish a separate instance of a directory server with different values. If you do not want those values to be exercised by all users, set up a new configuration set entry and run a separate server instance pointing to that configuration set entry for groups with special needs.

Figure 5–1 shows three separate directory server instances, each with a different value.

***Figure 5–1 Directory Entry Hierarchy Showing Multiple Configuration Set Entries***

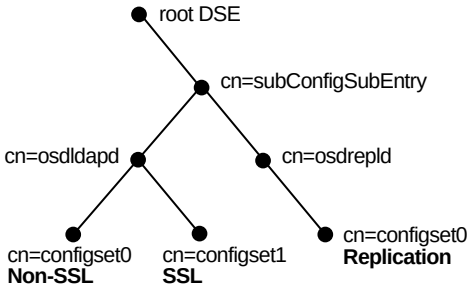


Figure 5–1 shows:

- An Oracle directory server (`cn=osdldap`) with:
  - One instance listening on the default port and using `configset0` with SSL set to *off*
  - A second instance listening on the SSL port and using `configset1` with SSL set to *on*
- A replication server instance (`cn=osdrepld`) using `configset0`

**See Also:**

- Chapter 8, "Managing Secure Sockets Layer (SSL)" for information about configuration parameters for SSL
- Chapter 10, "Managing Directory Replication" for information about configuration parameters for replication
- "Configuration Set Entry Attributes" on page E-5 for a list and descriptions of the entire set of attributes that are used to configure an instance of a directory server

## Managing Server Configuration Set Entries by Using Oracle Directory Manager

You can use Oracle Directory Manager to view, add, modify, and delete configuration set entries.

---

**Important Note:** You cannot change the parameters for an active instance directly; you must change the parameters in a configuration set entry and save it. After the configuration set entry is saved, use the OID Control Utility restart command to stop current Oracle directory server instances and restart them.

You can change a configuration set entry and start fresh instances that use the new parameters. The changes will not affect the older instances that are still running, however, unless they have been restarted.

For information on restarting directory server instances, see "Task 3: Reset the Default Security Configuration" on page 3-9.

---

### Viewing Configuration Set Entries by Using Oracle Directory Manager

To view configuration set entries:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server instance* > Server Management, then select Directory Server or Replication Server. The parameters of the active instance appear in the right pane.
2. Click a specific instance in the right pane. A Server Process dialog box appears.

   You can see all the parameters for the instance by selecting the tabs across the top of the dialog box. However, you cannot change them in this dialog box. To change them, you must change the configuration set entry on which they are based.

   **See Also:** "Modifying Configuration Set Entries by Using Oracle Directory Manager" on page 5-8

### Adding Configuration Set Entries by Using Oracle Directory Manager

The first time you add a configuration set entry, you can:

- Use the default configuration set as a template, then copy from the ones you create to make subsequent configuration sets
- Add a configuration set entry without copying from an existing one

**Adding a Configuration Set Entry by Copying from the Default Configuration Set** To add configuration set entries by copying the default configuration set entry:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server instance* > Server Management > Directory Server, then select Default Configuration Set.
2. On the toolbar, click the Create Like button. The Configuration Sets dialog box displays the General tab.
3. Fill in the fields with the information described in the following table:

| Field | Description |
|---|---|
| Max. Number of DB Connections | Type the number of concurrent database connections a single directory server process can have. The default is ten. |
| Number of Child Processes | Type the number of server processes a single instance can spawn. The default is one. |
| Set | Type the number of the configuration set entry. The default configuration set is 0. There can be as many different configuration sets as needed. The same configuration set can be used by more than one instance if the parameter needs of the multiple instances are the same. The set number is not modifiable. |

4. Select the SSL Settings tab and fill in the fields with the information described in this table:

| Field | Description |
|---|---|
| SSL Enable | Select to enable SSL authentication. If you do not select this check box, SSL is not enabled, and you do not need to set any other parameters on this page. |
| SSL Authentication | Choose one of the following:<br>■ No SSL Authentication—Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, SSL encryption / decryption only is used.<br>■ SSL Client and Server Authentication—Both client and server authenticate themselves to each other and send certificates to each other.<br>■ SSL Server Authentication—Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |

| Field | Description |
|---|---|
| SSL Wallet URL | Type the location of the SSL wallet. If you elect to change the location of the Oracle wallet, you must change this parameter. You must set the wallet location on both the client and the server. For example, on Solaris, you could set this parameter as follows:<br>`orclsslwalleturl=file:/Home/my_dir/my_wallet`<br>On Windows NT, you could set this parameter as follows:<br>`file:C:\my_dir\my_wallet` |
| SSL Wallet Password | Type the password for the server-side wallet. This password was set during creation of the wallet. If you change the password, you must change this parameter. |
| SSL Wallet Confirm Password | Retype the new password in this field when you change the password. |
| SSL Port | The default SSL port is 636. You can change the SSL port. |

**See Also:** Appendix C for information about setting the location of the Oracle Wallet and the Oracle Wallet password.

5. Click Apply.

**Note:** Remember: The changes will not affect the active directory server instance until you restart it. See "Restarting Directory Server Instances" on page 3-7.

**See Also:** "Setting Debug Logging Levels by Using the OID Control Utility" on page 5-23

**Adding a Configuration Set Entry Without Copying from an Existing Configuration Set** To create a new configuration set entry without copying from a previous configuration set entry:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server instance* > Server Management > Directory Server, then select Default Configuration Set.
2. On the toolbar, click Create. A Configuration Sets dialog box displays the General tab page. Fill in the fields as described in this table:

| Field | Description |
|---|---|
| Max. Number of DB Connections | Type the number of concurrent database connections a single directory server process can have. The default is ten. |
| Number of Child Processes | Type the number of server processes a single instance can spawn. The default is one. |
| Set | Type the number of the configuration set entry. The default configuration set is 0. There can be as many different configuration sets as needed. The same configuration set can be used by more than one instance if the parameter needs of the multiple instances are the same. The set number is not modifiable. |

3. Select the SSL Settings tab and fill in the fields with the information described in this table

| Field | Description |
|---|---|
| SSL Enable | Select to enable SSL authentication. If you do not select this check box, SSL is not enabled, and you do not need to set any other parameters on this page. |
| SSL Authentication | Choose one of the following:<br>■ No SSL Authentication—Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, SSL encryption / decryption only is used.<br>■ SSL Client and Server Authentication—Both client and server authenticate themselves to each other and send certificates to each other.<br>■ SSL Server Authentication—Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |

| Field | Description |
|---|---|
| SSL Wallet URL | Type the location of the SSL wallet. If you elect to change the location of the Oracle wallet, you must change this parameter. You must set the wallet location on both the client and the server. For example, on Solaris, you could set this parameter as follows:<br>`orclsslwalleturl=file:/Home/my_dir/my_wallet`<br>On Windows NT, you could set this parameter as follows:<br>`file:C:\my_dir\my_wallet` |
| SSL Wallet Password | Type the password for the server-side wallet. This password was set during creation of the wallet. If you change the password, you must change this parameter. |
| SSL Wallet Confirm Password | Retype the new password in this field when you change the password. |
| SSL Port | The default SSL port is 636. You can change the SSL port. |

Click Ok.

### Modifying Configuration Set Entries by Using Oracle Directory Manager

To modify configuration set entries:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server instance* > Server Management > Directory Server, then select the configuration set entry you want to modify. The configuration set appears in the group of tab pages in the right pane.

   Modify the values in the fields for the General tab as described in this table:

| Field | Description |
|---|---|
| Max. Number of DB Connections | Type the number of concurrent database connections a single directory server process can have. The default is ten. |
| Number of Child Processes | Type the number of server processes a single instance can spawn. The default is one. |
| Set | Type the number of the configuration set entry. The default configuration set is 0. There can be as many different configuration sets as needed. The same configuration set can be used by more than one instance if the parameter needs of the multiple instances are the same. The set number is not modifiable. |

You can change any of the values. Press Apply to save the changes.

2. Select the SSL Settings tab. Modify the fields as described in the following table.

| Field | Description |
|---|---|
| SSL Enable | Select to enable SSL authentication. If you do not select this check box, SSL is not enabled, and you do not need to set any other parameters on this page. |
| SSL Authentication | Choose one of the following:<br>■ No SSL Authentication—Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, SSL encryption / decryption only is used.<br>■ SSL Client and Server Authentication—Both client and server authenticate themselves to each other and send certificates to each other.<br>■ SSL Server Authentication—Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |
| SSL Wallet URL | Type the location of the SSL wallet. If you elect to change the location of the Oracle wallet, you must change this parameter. You must set the wallet location on both the client and the server. For example, on Solaris, you could set this parameter as follows:<br>`orclsslwalleturl=file:/Home/my_dir/my_wallet`<br>On Windows NT, you could set this parameter as follows:<br>`file:C:\my_dir\my_wallet` |
| SSL Wallet Password | Type the password for the server-side wallet. This password was set during creation of the wallet. If you change the password, you must change this parameter. |
| SSL Wallet Confirm Password | Retype the new password in this field when you change the password. |
| SSL Port | The default SSL port is 636. You can change the SSL port. |

3. Once you are satisfied with the parameters you have set for the new configuration set entry, click Apply.
4. Restart the server instance for the command to take effect.

**Note:** Remember: The changes will not affect the active directory server instance until you restart it. See "Restarting Directory Server Instances" on page 3-7.

**See Also:** Appendix C for information on setting the location of the Oracle Wallet and the Oracle Wallet password.

### Deleting Configuration Set Entries by Using Oracle Directory Manager

To delete configuration set entries:

1. In the navigator pane, expand Server Management > Directory Server.
2. In the navigator pane, select the configuration set entry you want to delete.
3. Click Delete on the toolbar.

**Note:** Remember: The changes will not affect the active directory server instance until you restart it. See "Restarting Directory Server Instances" on page 3-7.

## Managing Server Configuration Set Entries by Using Command Line Tools

Although changing configuration set entries by using Oracle Directory Manager is desirable, it can sometimes be more convenient to use the available command line tools—for example, when you want to make the same set of changes across multiple Oracle directory servers.

When you add or modify configuration set entries by using the command line tools, the input file for adding a new configuration set entry should be written in **LDAP Data Interchange Format (LDIF)**. It should contain only the attributes and values that differ from the installed defaults. The directory server uses the attribute values that you establish in the new configuration set entry to override its own existing values for these attributes.

**See Also:** "LDAP Data Interchange Format (LDIF) Syntax" on page A-2 for information on LDIF

### Adding Configuration Set Entries by Using ldapadd

If you are adding a new Oracle directory server instance, you can either use an existing configuration set entry, or add a new one for the new instance.

To add a new configuration set entry, create an input file, and then load the input file with ldapadd. Follow these steps:

1. Create the input file in a text editor.

   Input files must use LDIF format. When you create the input file, you need to define or include only those attributes that differ from the current values in that configuration set entry.

   In this example, the parameter `configset2` is the RDN, or local name, of the new entry, the wallet location is: `/HOME/test/wallet`, and the password is `welcome`.

   ```
   dn:cn=configset2, cn=oidldapd, cn=subconfigsubentry
   cn:configset2
   objectclass:orclConfigSet
   objectclass:orclLDAPSubConfig
   objectclass:top
   orclsslauthentication:1
   orclsslenable:1
   orclsslport:5000
   orclsslversion:3
   orclsslwalletpasswd:welcome
   orclsslwalleturl:file:/HOME/test/wallet
   ```

2. Run ldapadd with an input file.

   At the system prompt, type the command to add the input file. If the example shown above were given the file name `newconfigs`, the ldapadd command would look something like this:

   ```
   ldapadd [options] -f newconfigs
   ```

   **See Also:**

   - "LDAP Data Interchange Format (LDIF) Syntax" on page A-2
   - "ldapadd Syntax" on page A-4 for a detailed list of options available with this command
   - "Configuration Set Entry Attributes" on page E-5 for a description of configuration set entry attributes

### Modifying and Deleting Configuration Set Entries by Using ldapmodify

To modify or delete an existing configuration set entry, create an input file containing only the attributes that you want to change, and then load the input file with the ldapmodify command. Follow these steps:

1. Create the input file.

   When you create the input file, define or include only those attributes that differ from the installed defaults.

   Input files must have LDIF format.

   In the example shown below, the parameter `cn=configset2,cn=osdldapd,cn=subconfigsubentry` is the DN, or local name, of an existing configuration set entry. This example shows how to modify the ORCLSSLPORT parameter to 7000.

   ```
   dn:cn=configset2,cn=osdldapd,cn=subconfigsubentry
   changetype: modify
   replace: orclsslport
   orclsslport: 7000
   ```

2. Run ldapmodify referencing the input file.

   Type the command to reference the input file at the system prompt. For example, if the input file were named `configfile`, your ldapmodify command would look something like the command shown that follows:

   ```
   ldapmodify [options] -f configfile
   ```

   **See Also:**

   - "LDAP Data Interchange Format (LDIF) Syntax" on page A-2
   - "ldapmodify Syntax" on page A-13 for a more detailed discussion of ldapmodify, and a list of its options
   - "Configuration Set Entry Attributes" on page E-5 for a description of configuration set entry attributes

# Setting System Operational Attributes

Operational **attributes**—as opposed to application attributes—pertain to the operation of the directory itself. Some operational information is specified by the directory to control the server—for example, the time stamp for an entry. Other operational information, such as access information, is defined by administrators and is used by the directory program in its processing. You must have superuser privileges to set system operational attributes.

This section contains these topics:

- Setting System Operational Attributes by Using Oracle Directory Manager
- Setting System Operational Attributes by Using ldapmodify

**See Also:** "Kinds of Attribute Information" on page 2-5

## Setting System Operational Attributes by Using Oracle Directory Manager

You can view and set some of the operational attributes for each Oracle directory server to which you are connected by using **Oracle Directory Manager**. To do this, in the navigator pane, expand Oracle Internet Directory Servers, then select a server. System operational attributes appear in the right pane.

The next table describes the fields displayed in Oracle Directory Manager for each system operational attribute.

| Field | Description | Default Value | Modifiable? |
|---|---|---|---|
| Configuration Set Location | DN of the entry holding the top of the naming context in this server | `cn=subconfigsubentry` | No |
| Indexed Attribute Locations | DN for the file containing all indexed attributes | `cn=catalogs` | No |
| Naming Contexts | DN for the naming contexts contained in this server. Enter a new value in the field. If you are not sure of the value, click Browse to bring up a search window. | none | Yes |
| Oracle Directory Version | OID version / release that you are using | 2.1.1.0.0 | No |
| Password Encryption | Hash algorithm for encrypting the password. Options are:<br>■ **MD4**<br>■ **MD5**<br>■ No encryption<br>■ **SHA**<br>■ **UNIX Crypt** | MD4 | Yes |
| Process Instance Location | DN of the entry holding the Instance Registry in this server | `cn=subschemasubentry` | No |
| Query Entry Return Limit | Maximum number of entries to be returned by a search | 1000 | Yes |
| Replication Agreements | DN of the entry holding the replication agreement | `cn=orclareplagreements` | No |
| Replication Log Location | DN of the entry holding the change log in this server | `cn=changelog` | No |

| Field | Description | Default Value | Modifiable? |
|---|---|---|---|
| Replication Status Location | DN of the entry holding the change status in this server | `cn=changestatus` | No |
| Schema Definition Location | DN of the schema | `cn=subschemasubentry` | No |
| Server Mode | Determines whether data can be written to the server. Change the default to Read Only during replication process. | Read / Write | Choices are Read / Write and Read-Only |
| Server Operation Time Limit | Maximum amount of time, in seconds, allowed for a search to be completed | 3600 | Yes |

## Setting System Operational Attributes by Using ldapmodify

The modifiable system operational attributes are:

| Attribute | Description | Default |
|---|---|---|
| `namingContexts` | Topmost DNs for the naming contexts contained in this server. You must have super user privileges to publish a DN as a naming context. | none |
| `orclCryptoScheme` | Hash algorithm for encrypting the password. Options are:<br>■ MD4<br>■ MD5<br>■ No encryption<br>■ SHA<br>■ UNIX Crypt | MD4 |
| `orclSizeLimit` | Maximum number of entries to be returned by a search | 1000 |
| `orclServerMode` | Determines whether data can be written to the server. Change the default to Read-Only during replication process. | Read / Write |
| `orclTimeLimit` | Maximum amount of time, in seconds, allowed for a search to be completed | 3600 |

**See Also:** "ldapmodify Syntax" on page A-13 for a more detailed discussion of ldapmodify, and a list of its options

## Managing Naming Contexts

To enable users to search for specific naming contexts, you can publish those naming contexts. To do this, you specify the topmost entry of each naming context as a value of the `namingContexts` attribute in the root DSE.

For example, suppose you have a DIT with three major naming contexts, the topmost entries of which are `c=uk`, `c=us`, and `c=de`. If these entries are specified as values in the `namingContexts` attribute, then a user, by specifying the appropriate filter, can find information about them by searching the root DSE. The user can then focus the search—for example, by concentrating on the `c=de` naming context in particular.

To publish a naming context, you can use either Oracle Directory Manager or ldapmodify. The `namingContexts` attribute is multi-valued, so you can specify multiple naming contexts.

To search for published naming contexts, perform a base search on the root DSE with `objectClass` =* specified as a search filter. The retrieved information includes those entries specified in the `namingContexts` attribute.

Before you publish a naming context, be sure that:

- You are a directory administrator with the necessary access to the root DSE
- The topmost entry of that naming context exists in the directory

This section contains these topics:

### Publishing Naming Contexts by Using Oracle Directory Manager

1. In the navigator pane, expand Oracle Internet Directory Servers and select the directory server on which you want to specify a naming context. The corresponding tab pages for that directory server appear in the right pane.
2. In the System Operational Attributes tab page, in the Naming Contexts field, enter the topmost DN of the naming context you want to publish. You can also click Browse to open a search window.
3. Click Apply.

### Publishing Naming Contexts by Using ldapmodify

The following example input file specifies the entry `c=uk` as a naming context.

```
dn:
changetype: modify
add: namingcontexts
namingcontexts: c=uk
```

## Managing Password Encryption

During installation, you were prompted to set the encryption scheme for passwords. You can change that initial configuration by using either Oracle Directory Manager or ldapmodify. You must be a superuser to change the type of password encryption. This section contains these topics:

- Managing Password Encryption by Using Oracle Directory Manager
- Managing Password Encryption by Using ldapmodify

### Managing Password Encryption by Using Oracle Directory Manager

To change the type of password encryption by using Oracle Directory Manager:

1. In the navigator pane, expand Oracle Internet Directory Servers and select the directory server instance for which you want to reset password encryption. The corresponding tab pages for that directory server appear in the right pane.
2. In the System Operational Attributes tab page, in the Password Encryption field, select the type of password encryption you want to use. Options are:

- **MD4**
- **MD5**

- No encryption
- **SHA**
- **UNIX Crypt**

3. Click Apply.

### Managing Password Encryption by Using ldapmodify

The following example changes the password encryption algorithm to SHA:

```
ldapmodify -h myhost -p 389 -v <<EOF
dn:
changetype: modify
replace: orclcryptoscheme
orclcryptoscheme: SHA
EOF
```

**See Also:** "Password Encryption" on page 2-17

## Configuring Searches

You can set the maximum number of entries returned in searches, as well as the maximum amount of time, in seconds, for searches to be completed. You can do both of these by using either Oracle Directory Manager or ldapmodify.

This section contains these topics:

- Configuring Searches by Using Oracle Directory Manager
- Configuring Searches by Using ldapmodify

## Configuring Searches by Using Oracle Directory Manager

You can use Oracle Directory Manager to set the maximum number of retries returned in searches and the maximum amount of time to allow for searches.

### Setting the Maximum Number of Entries Returned in Searches by Using Oracle Directory Manager

1. In the navigator pane, expand Oracle Internet Directory Servers and select a directory server instance. The group of tab pages for that server appear in the right pane.
2. In the System Operational Attributes tab page, in the Query Entry Return Limit field, enter the maximum number of entries to be returned by a search. The default is 1000.
3. Click Apply.

### Setting the Maximum Amount of Time For Searches by Using Oracle Directory Manager

1. In the navigator pane, expand Oracle Internet Directory Servers and select a directory server instance. The group of tab pages for that server appear in the right pane.
2. In the System Operational Attributes tab page, in the Server Operation Time Limit, enter the maximum number of seconds for a search to be completed. The default is 3600.
3. Click Apply.

## Configuring Searches by Using ldapmodify

You can use ldamodify to set the maximum number of retries returned in searches and the maximum amount of time to allow for searches.

### Setting the Maximum Number of Entries Returned in Searches by Using ldapmodify

The following example changes the maximum number of entries to be returned in searches to 500.

```
ldapmodify -h myhost -p 389 -v <<EOF
dn:
changetype: modify
replace: orclsizelimit
orclsizelimit: 500
EOF
```

### Setting the Maximum Amount of Time For Searches by Using ldapmodify

The following example changes the maximum amount of time for a search to 2400.

```
ldapmodify -h myhost -p 389 -v <<EOF
dn:
changetype: modify
replace: orcltimelimit
orcltimelimit: 2400
EOF
```

**See Also:** "ldapmodify Syntax" on page A-13

# Managing Super, Guest, and Proxy Users

A **superuser** is a special directory administrator who typically has full access to directory information.

A **guest user** is one who is not an anonymous user, and, at the same time, does not have a specific user entry.

A **proxy user** is typically used in an environment with a middle tier such as a firewall. In such an environment, the end user authenticates to the middle tier. The middle tier then logs into the directory on the end user's behalf, but does so as a proxy user. A proxy user has the privilege to switch identities and, once it has logged into the directory, switches to the end user's identity. It then performs operations on the end user's behalf, using the **authorization** appropriate to that particular end user.

You can administer user names and passwords for the super, guest, and proxy users by using either Oracle Directory Manager or ldapmodify.

---

**Note:** It is possible to log on to the Oracle Directory Manager without giving a user name or password. If you do this, you have the privileges specified for an anonymous user. Anonymous users should have very limited privileges.

---

**See Also:** Chapter 9, "Managing Directory Access Control" for information on how to set access rights

This section contains these topics:

- Managing User Names and Passwords by Using Oracle Directory Manager
- Managing User Names and Passwords by Using ldapmodify

## Managing User Names and Passwords by Using Oracle Directory Manager

---

**Note:** The passwords for superusers, guest users, and proxy users are encrypted by default. You cannot modify them to send them in the clear.

---

To change a user name or password for a superuser, guest user, or a proxy user by using Oracle Directory Manager:

1. In the navigator pane, expand Oracle Internet Directory Servers.
2. Select a server. The group of tab pages for that server appear in the right pane.
3. Select the Passwords tab. This page displays the current user names and passwords for each type of user. Note that passwords are not displayed in the password fields.

The next table lists and describes the fields in the Passwords tab page.

| Field | Description |
|---|---|
| Super User Name | Type the super user name. The default is `cn=orcladmin`. |
| Super User Password | Type the super user password. The default is `welcome`. You should change this password immediately. |
| Guest Login Name | Type the guest login name. Guests have privileges determined by the **Access Control Policy Point**s (**ACPs**) in the directory. The default is `cn=guest`. |
| Guest Login Password | Type the guest login password. The default is `guest`. |
| Proxy Login Name | Type the proxy login name. Proxy users have privileges determined by the ACPs in the directory. The default is `cn=proxy`. |
| Proxy Login Password | Type the proxy login password. The default is `proxy`. |

4. Edit the appropriate field in the Password tab page. To save your changes, click Apply.

## Managing User Names and Passwords by Using ldapmodify

To change a user name or password for a superuser, a guest user, or a proxy user, use ldapmodify to modify these attributes:

| User Name/Password | Attribute |
|---|---|
| Super user name | `orclsuname` |
| Super user password | `orclsupassword` |
| Guest user name | `orclguname` |
| Guest user password | `orclgupassword` |
| Proxy user name | `orclprname` |
| Proxy user password | `orclprpassword` |

For example, to change the password of the super user to *superuserpassword*, use ldapmodify to modify the **DSE** by using an LDIF file containing the following:

```
dn:
changetype:modify
replace:orclsupassword
orclsupassword:superuserpassword
```

> **See Also:** "ldapmodify Syntax" on page A-13 for ldapmodify syntax and usage notes.

# Setting Debug Logging Levels

You can set debug logging levels by using either **Oracle Directory Manager** or the **OID Control Utility**.

This section contains these topics:

- Setting Debug Logging Levels by Using Oracle Directory Manager
- Setting Debug Logging Levels by Using the OID Control Utility

## Setting Debug Logging Levels by Using Oracle Directory Manager

1. In the navigator pane, expand Oracle Internet Directory Servers and select a server. The group of tab pages for that server appear in the right pane.
2. Select the Debug Flags tab.

   Ordinarily, you can leave the check boxes on this tab page unselected. However, to generate a log for a specific problem, use this tab page to specify the debug logging level.

## Setting Debug Logging Levels by Using the OID Control Utility

To set debug logging levels by using the OID Control Utility, restart the Oracle directory server using the `-debug` option for an LDAP server, and the `-d` flag for the replication server. Use the debug level number based on Table 5–1.

Because debug levels are additive, you need to sum together the numbers representing the functions that you want to activate, and use that sum in the command line option.

By default, debug logging is turned off. To turn it on, modify the **DSE** attribute `orcldebugflag` to the level you want. You can configure debug levels to one of the following levels.

To see debug log files generated by the OID Control Utility, navigate to $*ORACLE_HOME*/ldap/log.

Table 5–1 provides the complete list of debug logging levels.

***Table 5–1 Debug Logging Levels***

| Logging Level Value | Function |
|---|---|
| 1 | Trace function calls |
| 2 | Debug packet handling |
| 4 | Heavy trace debugging |
| 8 | Connection management |
| 16 | Print out packets sent and received |
| 32 | Search filter processing |
| 64 | Configuration file processing |
| 128 | Access control list processing |
| 256 | Stats log connections / operations / results |
| 512 | Stats log entries sent |
| 1024 | Print communication with the back-end |
| 2048 | Print entry parsing debugging |
| 4096 | Schema-related debugging |
| 32768 | Replication-specific debugging |
| 65535 | Enable all debugging |

For example, to trace function calls (1) and active connection management (8), enter 9 as the debug level (8 + 1 = 9) as follows:

```
oidctl server=oidldapd instance=1 flags='-debug 9' restart
oidctl server=oidrepld instance=1 flags='-h my_host -p 389 -d 9' restart
```

This example restarts both the Oracle directory server as well as the Oracle directory replication server with the debugging flags.

# Using Audit Log

The audit log records critical events on the Oracle directory server that are important from both a security and an operational point of view. An administrator can query the audit log using ldapsearch commands. Because the log generation is contingent upon events occurring on the server, only the Oracle directory server itself can create the log entries. You cannot add audit log entries with either the **Oracle Directory Manager** or the command line tools. Only the server can add entries.

The audit log is made up of regular directory entries, one entry for each event. You can specify search criteria using ldapsearch, and you can view the audit log entries by using Oracle Directory Manager.

By default audit logging is turned off. To turn it on, modify the **DSE** attribute `orclauditlevel` to the level you want. You can configure audit levels to audit selected events only.

**See Also:**

- "Auditable Events" on page 5-27 for a listing of audit levels
- "Searching for Audit Log Entries by Using Oracle Directory Manager" on page 7-5
- "Searching for Audit Log Entries by Using ldapsearch" on page 5-30
- "bulkdelete Syntax" on page A-22

This section contains these topics:

## Structure of Audit Log Entries

Each audit log entry contains the `orclAuditoc` **object class**. Like all other structural object classes, `orclAuditoc` inherits from `top`. Its attributes include:

| Attribute | Description |
|---|---|
| `orclsequence` | Used to create the name of the entry. The name is generated using a database sequence. |
| `orcleventtype` | Specifies the type of event that occurred. This is a catalogued attribute. |
| `orcleventtime` | Specifies the time at which the event occurred. This is formatted in **UTC (Coordinated Universal Time)**. UTC is indicated by a z at the end of the value. For example, `orcleventtime: 199811281010z` |
| `orcluserdn` | Specifies the identity of the user who logged into the Oracle directory server to perform the operation. This attribute is catalogued. |
| `orclopresult` | Specifies the outcome of the operation. It states either SUCCESS if the operation succeeds, or the reason why the operation failed. |
| `orclauditmessage` | Specifies the textual message. This attribute is not catalogued. |
| `objectclass` | Contains the preset values `top` and `orclauditoc`. |

Note that the audit log entries do not become part of a regular search result set even though the search filter can satisfy the query criteria. For example, a search with the condition `objectclass=top` does not yield results from the auditlog entries. Only a search with `cn=auditlog` as the base of the search can find audit log entries.

---

**Note:** By default, the attributes `orcleventtype` and `orcluserdn` are indexed at installation of Oracle Internet Directory. If you drop the indexes from these attributes, you cannot search for them. To re-create the index for these attributes, use the Catalog Management tool. See "Indexing an Attribute by Using Command Line Tools" on page 6-29.

---

**See Also:**

- "Catalog Management Tool Syntax" on page A-28 for information about catalogued attributes
- "Object Class Types" on page 2-9 for a description of `top`

## Position of Audit Log Entries in the DIT

The audit log container is part of the DSE. It holds its entries as children, organized according to the `orclsequence` attribute. See Figure 5–2.

***Figure 5–2 Sample Audit Log in DSE***

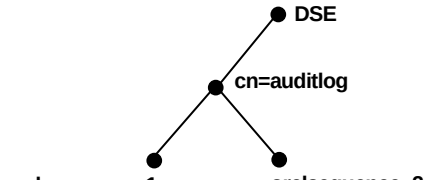


```
orclsequence: 1
orcleventtime: 199811281010z
orclauditmessage: Adding Attribute: (1.2.32.43.3.NAME 'myattr' SYNTAX '1.2.3.4.5.6.7')
orcleopresult: Invalid syntax.
orcluserdn: cn=orcladmin
objectclass: top
objectclass: orclauditoc
```

## Auditable Events

The next table shows the auditable events and their audit levels. The third column, Audit Levels, contains hexidecimal values. You can audit more than one event by adding their corresponding values found in this column.

| Event | Description | Audit Levels |
|---|---|---|
| Superuser login | Super user bind to the server (successes or failures) | 0x0001 |
| Schema element add/replace | Adding a new schema element (success and failure) | 0x0002 |
| Schema element delete | Deleting a schema (successes or failures) | 0x0004 |
| Bind | Unsuccessful bind cases | 0x0008 |

| Event | Description | Audit Levels |
|---|---|---|
| Access violation | Access denied by **ACP** | 0x0010 |
| DSE modification | Changes to DSE entry (successes or failures) | 0x0020 |
| Replication login | Replication server authentication (successes or failures) | 0x0040 |
| **ACL** modification | Changes to ACPs | 0x0080 |
| User password modification | Modification of user password attribute | 0x0100 |
| Add | ldapadd operation (successes or failures) | 0x0200 |
| Delete | ldapdelete operation (successes or failures) | 0x0400 |
| Modify | ldapmodify operation (successes or failures) | 0x0800 |
| ModifyDN | ldapModifyDN operation (successes or failures) | 0x1000 |

## Setting the Audit Level

Events described in the previous section can be turned on or off. The DSE attribute `orclauditlevel` indicates the current audit level set on the server. A value of 0 for the attribute means no auditing, which is the default.

You can set the audit level by using either Oracle Directory Manager or ldapmodify. Both methods are described in this section.

### Setting the Audit Level by Using Oracle Directory Manager

To set the audit level by using Oracle Directory Manager:

1. In the navigator pane, expand Oracle Internet Directory Servers and select the directory server instance.
2. In the right pane, select the Audit Mask Levels tab page.
3. Select the check box for the audit level you want to use.
4. Click Apply.

**Note:** Remember: The changes will not affect the active directory server instance until you restart it. See "Restarting Directory Server Instances" on page 3-7.

**See Also:** "Auditable Events" on page 5-27 for a description of each audit level

### Setting the Audit Level by Using ldapmodify

To audit more than one event, add the values of their the audit masks. For example, suppose you want to audit the following three events:

| Event | Audit Level | Value |
|---|---|---|
| Schema element delete | 0x0004 | 4 |
| DSE modification | 0x0020 | 32 |
| Add | 0x0200 | 512 |
| Total | | 548 |

The total value of the audit levels is 548. The ldapmodify command would therefore look something like this:

```
ldapmodify -p port -h host << EOF
dn:
changetype:modify
replace: orclauditlevel
orclauditlevel: 548
EOF
```

Restart the directory server instance after any changes are made to `orclauditlevel` for the changes to take effect.

**See Also:** "Task 3: Reset the Default Security Configuration" on page 3-9

## Searching for Audit Log Entries

You can search for audit log entries by using either Oracle Directory Manager or ldapsearch.

### Searching for Audit Log Entries by Using Oracle Directory Manager

**See:** "Searching for Audit Log Entries by Using Oracle Directory Manager" on page 7-5

### Searching for Audit Log Entries by Using ldapsearch

The **DN** for the audit log container is `cn=auditlog`. To search for audit log entries, perform a subtree or one-level search, with the container object `cn=auditlog` as the base of the search.

**See:** "ldapsearch Syntax" on page A-18

## Purging the Audit Log

You can use bulkdelete to purge audit log objects under the container `cn=auditlog`. Run the following command:

```
bulkdelete.sh -connect net_service_name -base "cn=auditlog"
```

# Viewing Active Server Instance Information

You can use **Oracle Directory Manager** to view information about any active server instance. To do this:

1. In the navigator pane, expand Oracle Internet Directory Servers and select a server. The group of tab pages for that server instance appear in the right pane.
2. Select the Server Management tab to display basic information—namely, type, instance number, debug level, and host name—for all active server instances.
3. To see configuration parameters for a particular server instance, select the server.
4. Click View Properties. The Server Process dialog box displays configuration parameters for the server instance you selected. Note that you cannot change configuration parameters in this dialog box. To change them, you must change the configuration set entry on which they are based.

**See Also:** "Managing Server Configuration Set Entries by Using Oracle Directory Manager" on page 5-4 for instructions on changing configuration set entries

## Changing the Password to an Oracle Data Server

The Oracle Internet Directory uses a password when connecting to an Oracle database. The default for this password when you install Oracle Internet Directory is ODS. You can change this password by using the **OID Database Password Utility**.

**See Also:** "OID Database Password Utility Syntax" on page A-37

# 6

# Managing the Directory Schema

This chapter explains how to administer the Oracle Internet Directory object classes and attributes.

This chapter contains these topics:

# About the Directory Schema

A directory schema does the following:

- Contains rules about the kinds of objects you can store in the directory
- Contains rules for how directory servers and clients treat information during operations such as a search
- Helps to maintain the integrity and quality of the data stored in the directory
- Reduces duplication of data
- Provides a predictable way for directory-enabled applications to access and modify directory objects

The directory schema contains all information about how data is organized in the DIT. It includes attribute types, and the syntaxes and matching rules that apply to them. It also contains the various groupings of attributes, called object classes.

This chapter discusses each of these elements.

**See Also:** "The Directory Schema" on page 2-12

# About Object Class Management

This section explains how to add and modify **object classes**. Oracle Corporation recommends that you understand the basic concepts of directory components before attempting to add to or modify the base schema in the directory.

**See Also:**

- "Object Classes" on page 2-8 for a conceptual overview of object classes
- Appendix E, "Schema Elements" for a list of schema components installed with Oracle Internet Directory

This section contains these topics:

- Guidelines for Adding Object Classes
- Guidelines for Modifying Object Classes
- Guidelines for Deleting Object Classes

## Guidelines for Adding Object Classes

When you add directory entries, you select object classes for those entries. The attributes of an entry are determined by the object classes to which that entry is assigned.

Entries must be loaded in a top-down sequence. When you add an entry, all of its parent entries must already exist in the directory. Similarly, when you add entries that reference object classes and attributes, those referenced object classes and attributes must already exist in the directory schema. In most cases this will not be a problem since the directory server is delivered with a full set of standard directory objects.

---

**Note:** Every schema object in the Oracle Internet Directory has certain limitations. For example, some objects cannot be changed. These limitations are explained as constraints and rules in this chapter.

---

The attributes that an entry **inherits** from an object class may be either mandatory or optional. Optional attributes need not be present in the directory entry.

You can specify for any object class whether an attribute is mandatory or optional; however, the characteristic you specify is binding only for that object class. If you place the attribute in another object class, you can again specify whether the attribute is mandatory or optional for that object class. You can:

- Select from existing standard object classes
- Add a new, non-standard object class and assign it existing attributes
- Modify an existing object class, assigning it a different set of attributes
- Add and modify existing attributes

**See Also:** "About Attribute Management" on page 6-16

Administrators typically assign object classes to entries based on the attributes present in that object class. However, **superclasses** let you take advantage of inheritance—that is, the object classes selected for an entry have a hierarchy of superclasses from which they inherit mandatory and optional attributes. By default, all object classes inherit from the `top` object class.

When you add or perform an operation on an entry, you do not need to specify the entire hierarchy of superclasses associated with that entry. This feature, called object

class explosion, enables you to specify only the leaf object classes. Oracle Internet Directory resolves the hierarchy for the leaf object classes and enforces the information model constraints. For example, the `inetOrgPerson` object class has `top`, `person` and `organizationalPerson` as its superclasses. When you create an entry for a person entry, you need to specify only `inetOrgPerson` as the object class. Oracle Internet Directory then enforces the schema constraints defined by the respective superclasses, namely, `top`, `person`, and `organizationalPerson`.

When you add object classes, keep the following guidelines in mind:

- Every structural object class must have `top` as a superclass.
- The name and the object identifier of an object class must be unique across all the schema components.
- Schema components referred to in the object class, such as superclasses, must already exist.
- The superclass of an abstract object class must be abstract also.
- It is possible to redefine mandatory attributes in a superclass into optional attributes in the new object class. Conversely, optional attributes in a superclass can be redefined into mandatory attributes in the new object class.

**See Also:** "Subclasses, Superclasses, and Inheritance" on page 2-9 for a conceptual discussion of these terms

## Guidelines for Modifying Object Classes

This section discusses the types of modifications you can make to an existing object class. You can perform modifications through Oracle Directory Manager and through the command line tools.

You can make these changes to an object class:

- Change a mandatory attribute into an optional attribute
- Add optional attributes
- Add additional superclasses
- Convert *abstract* object classes into *structural* or *auxiliary* object classes unless the abstract object class is a superclass to another abstract object class

When you modify object classes, keep these guidelines in mind:

- You cannot modify an object class that is part of the standard LDAP schema. You can, however, modify user-defined object classes. Also, if existing object

classes do not have the attributes you need, you can create an auxiliary object class and associate the needed attributes with it.

- You cannot add additional mandatory attributes to an existing object class.
- You cannot modify object classes in the base schema.
- You cannot remove attributes or superclasses from an existing object class.
- You cannot convert structural object classes to other object class types.
- You should not modify an object class if there are entries already associated with it.

**See Also:**

- "Managing Object Classes by Using Oracle Directory Manager" on page 6-6
- "Managing Object Classes by Using Command Line Tools" on page 6-14

## Guidelines for Deleting Object Classes

There are also some limitations on deleting object classes:

- You cannot delete object classes from the base schema.
- You can delete object classes that are not in the base schema as long as they are not directly or indirectly referenced by other schema components. For example, there may be some directory entries referring to these object classes. Deleting these object classes renders these entries inaccessible.

---

**Note:** Oracle Internet Directory does not enforce these rules. They are provided here as guidelines.

---

# Managing Object Classes by Using Oracle Directory Manager

This section contains these topics:



## Searching for Object Classes by Using Oracle Directory Manager

You can specify your search for an object class by:

- Selecting an object class property, for example, a name or an object identifier
- Entering a value for the property you selected
- Selecting a search filter specifying the relationship between the object class property you selected and the value you entered, for example, Begins With or Exactly Matches

This section provides more details on how to enter an object class search.

To search for an object class:

1. In the navigator pane, select Schema Management. The Schema Management tab pages appear in the right pane.
2. Click the Find Object Classes button at the lower right of the right pane, or, from the menu bar, click Edit > Find Object Classes. The Find: Object Classes dialog box appears.

3. In the menu farthest to the left on the search criteria bar, select the property of the object class for which you want to search. Options are:

| Option | Description |
|---|---|
| Name | Name of the object class for which you are searching. For example, the phrase `Name Exact Match subAcl` gives you the `subAcl` object class. |
| Object ID | Object Identifier for the object class for which you are searching. For example, the phrase `Object ID Begins With 2.5.2` gives you a list of object classes whose object identifiers begin with 2.5.2. |
| Description | Word in the description field. For example, the phrase `Description Contains Shoe` gives you a list of object classes with the word *shoe* in the description column. |
| Type | Type of object class for which you are searching, whether abstract, structural, or auxiliary |
| Superclass | Class from which the object class for which you are searching is derived |
| Mandatory Attributes | Mandatory attributes of the object class for which you are searching. For example, the phrase `Mandatory Attributes Contains cn` gives you a list of all object classes in which the `cn` attribute is mandatory. |
| Optional Attributes | Optional attributes of the object class for which you are searching |

**Note:** Not all attributes are used in every object class. Be sure that the attribute you specify actually corresponds to one in the object class for which you are looking. Otherwise, the search will fail.

4. In the text box at the right end of the search criteria bar, type the value of the property of the object class for which you are searching. For example, to search for all object classes in which the name of the object class begins with the letters `orcl`, type those letters in the text box at the right end of the search criteria bar.

5. In the menu in the middle of the search criteria bar, select the filter you want to use for your search. Options are:

| Filter | Description |
|---|---|
| Begins With | Searches by using only the first few characters of the property of the object class for which you are searching. For example, the phrase `Type Begins With aux` gives you a list of all of the auxiliary object classes. |
| Ends With | Searches by using only the last few characters of the property of the object class for which you are searching. For example, the phrase `Type Ends With ral` gives you a list of all of the structural object classes. |
| Contains | Searches for object classes in which the property you selected includes, but is not necessarily limited to, the value you enter. For example, the phrase `Optional Attributes Contains cn` gives you a list of all object classes in which `cn` is an optional attribute. |
| Exact Match | Searches for an object class in which the property you selected is exactly the same as the value you enter. For example, the phrase `Super Class Exact Match person` gives you a list of all object classes that have `person` as their superclass. |
| Greater Or Equal | Searches for an object class in which the property you selected is numerically or alphabetically greater than or equal to the value you enter. For example, the phrase `Name Greater or Equal orcl` gives you a list of object classes from those beginning with the letters `orcl` to those beginning with letters at the end of the alphabet. |
| Less or Equal | Searches for an object class in which the property you selected is numerically or alphabetically less than or equal to the value you enter. For example, the phrase `Name Less or Equal orcl` gives you a list of object classes from those beginning with the letters `orcl` to those at the beginning of the alphabet. |
| Not Null | Searches for all object classes in which the property you selected is present. For example, the phrase `Mandatory Attributes Not Null` gives you a list of all object classes which contain mandatory attributes. |

6. Below the Search Criteria field are five buttons described in the next table. Use these buttons to further refine your search.

| Button | Description |
|---|---|
| New | Creates a new search criteria bar in the Search Criteria field. This button is enabled only when the search criteria bar has been deleted. |
| And | Creates another search criteria bar in the Search Criteria field. Matches all object classes having one specified criterion with those that also have another specified criterion. |
| Or | Creates another search criteria bar in the Search Criteria field. Matches all object classes with either one specified attribute or another. |
| Not | Negates the criterion in the selected search criteria bar and retrieves all object classes that do not have the specified criterion. |
| Delete | Deletes a selected search criteria bar |

7. Click Search. The results of your search appear in the window at the lower portion of the Find:Object Class dialog box.

## Viewing Properties of Object Classes by Using Oracle Directory Manager

To view all object classes in the schema:

1. In the navigator pane, expand Schema Management. The tabs in the Schema Management pane display the components of the schema:
   - Object classes
   - Attributes
   - Syntaxes
   - Matching Rules
2. In the right pane, select the Object Classes tab page.

To examine an individual object class and its attributes, in the Object Classes tab page, click the object class. The properties of the selected object class appear in the Object Class dialog box.

3. In the Object Class dialog box:
   - Object classes from which attributes may be inherited are listed in the Super Class box
   - Mandatory attributes are listed in the Mandatory Attributes box
   - Optional attributes are listed in the Optional Attributes box

   Each box indicates whether the attributes are indexed so that they can be used in a search expression.

## Adding Object Classes by Using Oracle Directory Manager

To add object classes by using Oracle Directory Manager:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server*, then select Schema Management
2. Choose one of the following methods:
   - In the right pane, select the Object Classes tab and click the Create button in the toolbar.
   - Click the Create button at the bottom of the right pane.
   - From Operations menu, select Create Object Class.

   The New Object Class dialog box appears.

   Alternatively, select an object class that is similar to one you would like to create, and then click Create Like. A dialog box appears; it includes the attributes of the selected object class. You can create the new object class using the selected one as a template.

**3.** Enter the information in the fields described in the following table:

| Field | Description |
|---|---|
| Name | Enter the name of the object class you are creating. |
| Object ID | Enter the object identifier. This is a standardized numerical sequence based on IETF standards. It must be unique, and should comply with the system established within your organization. Normally it is derived from the identifier assigned by registration agencies, such as ANSI or ISO. |
| Description | Use this optional field for your information only. |
| Type | Specify the type of object class: Abstract, Structural, Auxiliary, None. |
| Super Class | Specify the class(es) from which to derive this object class. This object class will inherit all the attributes of the superclass(es) you select. Every structural object class must have `top` as one of its superclasses. Clicking Add displays the Super Class Selector dialog box from which you can select the superclass(es) you want to add. |
| Mandatory Attributes | Specify the attributes for which values must be entered. Clicking Add displays the Mandatory Attributes Selector dialog box from which you can select the mandatory attributes you want to add. |
| Optional Attributes | Specify the attributes for which values are not required. Clicking Add displays the Optional Attributes Selector dialog box from which you can select the optional attributes you want to add. |

**4.** Click OK.

**See Also:**

- "Object Class Types" on page 2-9
- "Subclasses, Superclasses, and Inheritance" on page 2-9
- Oracle Directory Manager online help for further details about adding object classes

## Modifying Object Classes by Using Oracle Directory Manager

To modify an object class:

1. In the navigator pane, select Schema Management, then select the Object Classes tab.
2. In the Object Classes tab page, double-click the object class you want to modify. The Object Class dialog box appears.
3. Modify or add the information in the fields described in the following table.

| Field | Description |
|---|---|
| Name | Enter the name of the object class you are creating. |
| Object ID | Enter the object identifier. This is a standardized numerical sequence based on IETF standards. It must be unique, and should comply with the system established within your organization. Normally it is derived from the identifier assigned by registration agencies, such as ANSI or ISO. |
| Description | Use this optional field for your information only. |
| Type | Specify the type of object class: Abstract, Structural, Auxiliary, None. |
| Super Class | Specify the class(es) from which to derive this object class. This object class will inherit all the attributes of the superclass(es) you select. Every structural object class must have `top` as one of its superclasses. Clicking Add displays the Super Class Selector dialog box from which you can select the superclass(es) you want to add. |
| Mandatory Attributes | Specify the attributes for which values must be entered. Clicking Add displays the Mandatory Attributes Selector dialog box from which you can select the mandatory attributes you want to add. |
| Optional Attributes | Specify the attributes for which values are not required. Clicking Add displays the Optional Attributes Selector dialog box from which you can select the optional attributes you want to add. |

4. Click OK.

**See Also:**

- "Object Class Types" on page 2-9
- "Subclasses, Superclasses, and Inheritance" on page 2-9

## Deleting Object Classes by Using Oracle Directory Manager

**Caution:** Oracle Corporation recommends that you not delete object classes from the schema.

Should you decide to delete an object class, be careful not to delete one that is in use or that you might want to use in the future. If you delete an object class that is referenced by any entries, those entries then become inaccessible.

**Note:** You can add attributes to an auxiliary object class or a user-defined structural object class.

**See Also:** Example: Adding a New Attribute to an Auxiliary or User-Defined Object Class on page 6-15 for an example of adding attributes to an auxiliary object class

To delete an object class by using Oracle Directory Manager:

1. In the navigator pane, select Schema Management.
2. In the right pane, select the Object Classes tab and select the object class you want to delete.
3. Click Delete.

# Managing Object Classes by Using Command Line Tools

You can use command line tools to add or modify existing object classes in the directory schema. The command line tools enable you to use input files. Furthermore, the commands can be batched together in scripts.

To add or modify schema components, use ldapmodify.

**See:** "ldapmodify Syntax" on page A-13

This section contains these examples:

- Example: Adding a New Object Class
- Example: Adding a New Attribute to an Auxiliary or User-Defined Object Class

## Example: Adding a New Object Class

To add a new object class schema component by using ldapmodify, at the system prompt type a command using the following syntax:

```
ldapmodify -h host -p port -f ldif_filename
```

For example:

```
ldapmodify –h myhost -p 389 -f new_object_class.ldi
```

In this example, the LDIF input file, `new_object_class.ldi`, contains data similar to this:

```
dn: cn=subschemasubentry
changetype: modify
add: objectclasses
objectclasses: ( 1.2.3.4.5 NAME 'myobjclass' SUP top STRUCTURAL MUST ( cn $
sn ) MAY ( telephonenumber $ givenname $ myattr ) )
```

The example above adds the *structural* object class named `myobjclass`, giving it an object identifier of `1.2.3.4.5`, specifying `top` as its superclass, requiring `cn` and `sn` as mandatory attributes, and allowing `telephonenumber`, `givenname`, and `myattr` as optional attributes. Note that all the attributes mentioned must exist prior to the execution of the command.

Be sure to leave the mandatory space between the opening and closing parentheses and the object identifier.

To create an *abstract* object class, follow the above example, replacing the word STRUCTURAL with the word ABSTRACT.

## Example: Adding a New Attribute to an Auxiliary or User-Defined Object Class

To add a new attribute to either an auxiliary object class or a user-defined structural object class, use ldapmodify. This example deletes the old object class definition and adds the new definition in a compound modify operation. The change is committed by the Oracle directory server in one transaction. Existing data is not affected. The input file should be as follows:

```
dn: cn=subschemasubentry
changetype: modify
delete: objectclasses
objectclasses: old value
-
add: objectclasses
objectclasses: new value
```

For example, to add the attribute changes to the existing object class country, the input file would be:

```
dn: cn=subschemasubentry
changetype: modify
delete: objectclasses
objectclasses:  ( 2.5.6.2 NAME 'country' SUP top STRUCTURAL MUST c MAY
( searchGuide $ description  )  )
-
add: objectclasses
objectclasses:  ( 2.5.6.2 NAME 'country' SUP top STRUCTURAL MUST c MAY
( searchGuide $ description  $ changes )  )
```

# About Attribute Management

This section contains these topics:

- Rules for Adding Attributes
- Rules for Modifying Attributes
- Rules for Deleting Attributes

You need to understand attributes from a conceptual standpoint before attempting operations involving attributes.

In most cases, the attributes available in the base schema will suit the needs of your organization. However, if you decide to use an attribute not available in the base schema, you can add a new attribute or modify an existing one.

By default, attributes are multi-valued. You can specify an attribute as single-valued by using either Oracle Directory Manager or command line tools.

> **See Also:** "Attributes" on page 2-3 for a conceptual discussion of attributes

## Rules for Adding Attributes

The rules for adding attributes are:

- The name and the object identifier of an attribute must be unique across all the schema components.
- Syntax and matching rules must agree.
- Any super attributes must already exist.

## Rules for Modifying Attributes

The rules for modifying attributes are:

- The name and the object identifier of an attribute must be unique across all the schema components.
- The syntax of an attribute cannot be modified.
- A single-valued attribute can be made into multi-valued, but a multi-valued attribute cannot be made single-valued.
- You cannot modify or delete base schema attributes.

### Rules for Deleting Attributes

The rules for deleting attributes are:

- Attributes from the base schema cannot be deleted.
- You can delete any attribute that is not referenced directly or indirectly by some other schema component.

  If you delete an attribute that is referenced by any entry, that entry will no longer be available for directory operations.

## Managing Attributes by Using Oracle Directory Manager

This section contains these topics:



### Searching for Attributes by Using Oracle Directory Manager

To search for attributes by using Oracle Directory Manager:

1. In the navigator pane, select Schema Management. The Schema Management tab pages appear in the right pane.
2. Select the Attributes tab page.
3. Click the Find Attributes button in the lower right corner. The Find Attributes dialog box appears

4. In the menu at the left end of the search criteria bar, select the property of the attributes for which you want to search. Options are:

| Field | Description |
|---|---|
| Name | Name of the attribute for which you are searching |
| Indexed | List of indexed attributes |
| Object ID | Object Identifier for the attribute for which you are searching. For example, the phrase `Object ID Begins With 2.5.2` gives you a list of attributes whose object identifiers begin with `2.5.2`. |
| Description | Words in the description column of attributes |
| Syntax | The standardized rules for data entry applicable to this attribute type. Use this to narrow your search to attributes using a particular syntax. |
| Size | Maximum size allowed for this object |
| Usage | Standards specifying how the attribute can be used. You narrow your search by entering one of the following options: `userApplications`, `directoryOperation`, `distributedOperation`, and `dSAOperation`. |
| Ordering | Standards specifying how precedence is established for values |
| Equality | Standards specifying how equality is determined in compare and search operations |
| Substring | Used for regular expression matching |
| Single Value | Indicates that this attribute type contains a maximum of one value |
| Super | Super attribute for the attribute for which you are searching |

5. In the text box at the right end of the search criteria bar, type part or all of the value of the attribute for which you want to search. For example, to search for all attributes whose names begin with the letters `orcl`, you would type those letters in the text box at the right end of the search criteria bar and create the phrase `Name Begins With orcl`.

**6.** In the menu in the middle of the search criteria bar, select the filter you want to use for your search. Options are:

| Option | Description |
|---|---|
| Begins With | Searches by using only the first few characters of the property's value. For example, the phrase `Syntax Begins With 1.3` gives you a list of all attributes in which the first few numbers of the syntax identifier are *1.3*. |
| Ends With | Searches by using only the last few characters of the property's value. For example, the phrase `Name Ends With License` gives you a list of all attributes with that ending, such as `carLicense`. |
| Contains | Searches for attributes that include the property with the value you enter. For example, the phrase `Ordering Contains time` gives you a list of all attributes with the word `time` in the Ordering column |
| Exact Match | Searches for a value that is exactly the same as that found in the attribute property you specified. For example, the phrase `Equality Exact Match caseIgnoreMatch` gives you a list of all attributes that have the `caseIgnoreMatch` matching rule. |
| Greater or Equal | Searches for an attribute that has a property that is numerically or alphabetically greater than or equal to the value you enter. For example, the phrase `Name Greater or Equal orcl` gives you a list of attributes from those beginning with `orcl` to those beginning with letters at the end of the alphabet. |
| Less or Equal | Searches for an attribute that has a property that is numerically or alphabetically less than or equal to the value you enter. For example, the phrase `Name Less or Equal orcl` gives you a list of attributes from those beginning with `orcl` to those beginning with letters at the start of the alphabet. |
| Not Null | Searches for all attributes in which the attribute property you selected is present. For example, the phrase `Description Not Null` gives you a list of all attributes which have text in the description field. |

7. Beneath the Search Criteria field are five buttons described in the following table. Use these buttons to further refine your search.

| Button | Description |
|---|---|
| New | Creates a new search criteria bar in the Search Criteria field. This button is enabled only when the Search Criteria field is empty. |
| And | Creates another search criteria bar in the Search Criteria field. Matches all attributes with one specified property with those that also have another specified property. |
| Or | Creates another search criteria bar in the Search Criteria field. Matches all attributes with either one specified property or another. |
| Not | Negates the criteria in the selected search criteria bar and matches all attributes that do not have the property specified. |
| Delete | Deletes a selected search criteria bar |

8. Click Search. The results of your search appear in the window at the lower portion of the Find: Attributes dialog box.

## Adding an Attribute by Using Oracle Directory Manager

This section contains these topics:

- Adding a New Attribute by Using Oracle Directory Manager
- Creating a New Attribute from an Existing One by Using Oracle Directory Manager

> **Tip:** Because equality, syntax, and matching rules are numerous and complex, it may be simpler to copy these characteristics from a similar existing attribute.

### Adding a New Attribute by Using Oracle Directory Manager

To add a new attribute:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server*, then select Schema Management.

2. Do one of the following:

   - In the right pane, select the Attributes tab, then click the Create button in the toolbar.
   - In the right pane, select the Attributes tab, then click the Create button at the bottom of the Attributes tab page.
   - From the Operation menu, select Create Attribute.

   The New Attribute Type dialog box appears. It contains two tab pages—General and Advanced—with fields in which you either enter values or select from menus.

3. In the General tab, enter values in each of the fields as described in the following table:

| Field | Description |
|---|---|
| Name | Type the name for this attribute. |
| Object ID | Type the Object ID for this attribute. The Object ID is a standardized numerical sequence based on IETF standards. It must be unique. Normally this is derived from the identifier assigned by registration agencies, such as ANSI or ISO.<br>For an explanation of the standard identifiers, see the current LDAP standards available through the IETF website. |
| Description | This optional field is for your information only. |
| Syntax | Type the standardized rules for data entry applicable to this attribute type. |
| Size | Type the maximum size allowed for this object. |
| Single Value | Select this check box to indicate that this attribute type contains a maximum of one value. |

4. Select the Advanced tab. Enter values in each of the fields as described in the following table.

| Field | Description |
|---|---|
| Indexed | Select to add this attribute to the index, thereby making it available for use in a search. Only those attributes that have an equality matching rule can be indexed. |
| Usage | Specify standards for how the attribute can be used. Options are:<br>■ `userApplications`<br>Attributes whose values must be entered by the user, for example, `telephoneNumber`<br>■ `directoryOperation`<br>Attributes whose values are entered by the directory server, for example, `creatorName` or `timeStamp`<br>■ `distributedOperation`<br>■ `dSAOperation`<br>Attributes used for the internal operation of the server, for example, orclUpdateSchedule |
| Ordering | Specify standards for how precedence is established for values |
| Equality | Specify standards for how equality is determined in compare and search operations |
| Substring | Specify regular expression matching. |
| Super | Add the super attribute for this attribute. To do this:<br>1. Click the Add button next to this field. The Super Attribute Selector appears.<br>2. Select the super attribute and click Select.<br>3. Repeat as needed.<br>To delete a super attribute from the Super field, select it, then click Delete. |

5. Click OK.

**Note:** To use this attribute, remember to declare it to be part of the attribute set for an object class. You do this by selecting Schema Management in the navigator pane, then, in the right pane, selecting the Object Classes tab page. For further instructions, see "Guidelines for Modifying Object Classes" on page 6-4.

### Creating a New Attribute from an Existing One by Using Oracle Directory Manager

To add an attribute by copying an existing attribute:

1. In the navigator pane, select Schema Management.
2. In the right pane, select the Attributes tab.
3. In the Attributes tab page, select the attribute you want to copy.
4. Click the Create Like button at the bottom of the right pane. The New Attribute Type dialog box for that attribute appears. This dialog box contains two tab pages—General and Advanced—with fields in which you enter values either by typing or selecting from menus.
5. Select the General tab and enter values in each of the fields as described in the following table. You must always change the DN to that of the new attribute.

| Field | Description |
|---|---|
| Name | Type the name for this attribute. |
| Object ID | Type the Object ID for this attribute. The Object ID is a standardized numerical sequence based on IETF standards. It must be unique. Normally this is derived from the identifier assigned by registration agencies, such as ANSI or ISO.<br>For an explanation of the standard identifiers, see the current LDAP standards available through the IETF website. |
| Description | This optional field is for your information only. |
| Syntax | Type the standardized rules for data entry applicable to this attribute type. |
| Size | Type the maximum size allowed for this object. |
| Single Value | Select this check box to indicate that this attribute type contains a maximum of one value. |

6. Select the Advanced tab and enter values in each of the fields as described in the following table.

| Field | Description |
|---|---|
| Indexed | Select to add this attribute to the index, thereby making it available for use in a search. Only those attributes that have an equality matching rule can be indexed. |
| Usage | Specify standards for how the attribute can be used. Options are:<br>■ `userApplications`<br>Attributes whose values must be entered by the user, for example, `telephoneNumber`<br>■ `directoryOperation`<br>Attributes whose values are entered by the directory server, for example, `creatorName` or `timeStamp`<br>■ `distributedOperation`<br>■ `dSAOperation`<br>Attributes used for the internal operation of the server, for example, orclUpdateSchedule |
| Ordering | Specify standards for how precedence is established for values |
| Equality | Specify standards for how equality is determined in compare and search operations |
| Substring | Specify regular expression matching. |
| Super | Add the super attribute for this attribute. To do this:<br>1. Click the Add button next to this field. The Super Attribute Selector appears.<br>2. Select the super attribute and click Select.<br>3. Repeat as needed.<br>To delete a super attribute from the Super field, select it, then click Delete. |

7. Click OK.

## Modifying an Attribute by Using Oracle Directory Manager

To modify an attribute by using Oracle Directory Manager:

1. In the navigator pane, select Schema Management.
2. In the right pane, select the Attributes tab and double-click an editable attribute in the list. The Attribute dialog box displays two tab pages—General and Advanced—with fields in which you enter values either by typing or selecting from menus.
3. Select the General tab and enter values in each of the fields as described in the following table.

| Field | Description |
|---|---|
| Name | Type the name for this attribute. |
| Object ID | Type the Object ID for this attribute. The Object ID is a standardized numerical sequence based on IETF standards. It must be unique. Normally this is derived from the identifier assigned by registration agencies, such as ANSI or ISO.<br>For an explanation of the standard identifiers, see the current LDAP standards available through the IETF website. |
| Description | This optional field is for your information only. |
| Syntax | Type the standardized rules for data entry applicable to this attribute type. |
| Size | Type the maximum size allowed for this object. |
| Single Value | Select this check box to indicate that this attribute type contains a maximum of one value. |

4. Select the Advanced tab and enter values in each of the fields as described in the following table.

| Field | Description |
|---|---|
| Indexed | Select to add this attribute to the index, thereby making it available for use in a search. Only those attributes that have an equality matching rule can be indexed. |
| Usage | Specify standards for how the attribute can be used. Options are:<br>■ `userApplications`<br>Attributes whose values must be entered by the user, for example, `telephoneNumber`<br>■ `directoryOperation`<br>Attributes whose values are entered by the directory server, for example, `creatorName` or `timeStamp`<br>■ `distributedOperation`<br>■ `dSAOperation`<br>Attributes used for the internal operation of the server, for example, orclUpdateSchedule |
| Ordering | Specify standards for how precedence is established for values |
| Equality | Specify standards for how equality is determined in compare and search operations |
| Substring | Specify regular expression matching. |
| Super | Add the super attribute for this attribute. To do this:<br>1. Click the Add button next to this field. The Super Attribute Selector appears.<br>2. Select the super attribute and click Select.<br>3. Repeat as needed.<br>To delete a super attribute from the Super field, select it, then click Delete. |

5. Click OK.

## Indexing an Attribute When You Create It

Oracle Internet Directory uses indexes to make attributes available for searches. When Oracle Internet Directory is installed, certain attributes are already indexed. If you want to use additional attributes in search filters, you must index them.

---

**Note:** You can use Oracle Directory Manager to index an attribute only at the time when you create it. You cannot use Oracle Directory Manager to index an already existing attribute. To index an already existing attribute, use the Catalog Management tool.

Also, only those attributes that have an equality matching rule can be indexed.

---

**See Also:** "Indexing an Attribute by Using Command Line Tools" on page 6-29 for instructions on using the command line catalog management tool

This section contains these topics:

- Viewing Indexed Attributes by Using Oracle Directory Manager
- Indexing an Attribute When You Create It by Using Oracle Directory Manager
- Dropping an Index from an Attribute by Using Oracle Directory Manager

### Viewing Indexed Attributes by Using Oracle Directory Manager

To view indexed attributes:

1. In the navigator pane, select Schema Management.
2. In the right pane, select the Attributes tab. The Attributes tab displays all of the attributes in the schema. A selected check box in the Indexed column indicates an indexed attribute.

### Indexing an Attribute When You Create It by Using Oracle Directory Manager

When you create an attribute as described in "Adding an Attribute by Using Oracle Directory Manager" on page 6-20, you use the New Attribute Type dialog box. On the Advanced tab page of that dialog box, you select the Indexed check box.

### Dropping an Index from an Attribute by Using Oracle Directory Manager

To drop an index from an attribute:

1. In the navigator pane, select Schema Management.
2. In the right pane, select the Attributes tab.
3. Select the indexed attribute. Note that this must be an attribute that is editable as indicated by the icon to the left of the attribute name.
4. Click Drop Index.

## Managing Attributes by Using Command Line Tools

This section discusses adding, modifying, and indexing attributes by using command line tools. This section contains these topics:

- Adding and Modifying Attributes by Using ldapmodify
- Indexing an Attribute by Using Command Line Tools

### Adding and Modifying Attributes by Using ldapmodify

> **See Also:** "ldapmodify Syntax" on page A-13 for a detailed explanation of this command and its options

To add a new attribute to the schema by using ldapmodify, type a command similar to the following at the system prompt:

```
ldapmodify -h host -p port -f ldif_filename
```

The LDIF file contains data similar to this:

```
dn: cn=subschemasubentry
changetype: modify
add: attributetypes
attributetypes: ( 1.2.3.4.5 NAME 'myattr' SYNTAX
                 '1.3.6.1.4.1.1466.115.121.1.38' )
```

To specify an attribute as single-valued, include in the attribute definition entry in the LDIF file the keyword SINGLE-VALUE with surrounding white space.

You can find a given syntax Object ID by using either Oracle Directory Manager or the ldapsearch command line tool.

**Viewing Syntaxes by Using Oracle Directory Manager** To view syntaxes by using Oracle Directory Manager:

1. In the navigator pane, select Schema Management.
2. In the right pane, select the Syntaxes tab.

**Viewing Syntaxes by Using by Using ldapsearch** Use ldapsearch on the subentry `cn=subSchemaSubentry`.

**See Also:** "ldapsearch Syntax" on page A-18

## Indexing an Attribute by Using Command Line Tools

This section discusses these topics:



### About Indexing

Oracle Internet Directory uses indexes to make attributes available for searches. When Oracle Internet Directory is installed, the entry `cn=catalogs` lists available attributes that can be used in a search.

If you want to use additional attributes in search filters, you must add them to the catalog entry. Only those attributes that have an equality matching rule can be indexed.

You can index a new attribute—that is, one for which no data exists in the directory—by using ldapmodify. You can index an attribute for which data already exists in the directory by using the Catalog Management tool. You can drop an index from an attribute by using ldapmodify, but the recommended method is by using the Catalog Management tool.

### Indexing an Attribute for Which *No* Directory Data Exists by Using ldapmodify

Once you have defined a new attribute in the schema, you can add it to the catalog entry by using ldapmodify.

To add an attribute for which no directory data exists by using ldapmodify, import an LDIF file by using ldapmodify. For example, to add a new attribute `foo` that has already been defined in the schema, import the following LDIF file by using ldapmodify:

```
dn: cn=catalogs
Changetype: modify
Add: orclindexedattribute
Orclindexedattribute: foo
```

You should not use this method to index an attribute for which data exists in the directory. To index such an attribute, use the Catalog Management Tool.

To drop an index from an attribute by using ldapmodify, specify `delete` in the LDIF file. For example:

```
dn: cn=catalogs
Changetype: modify
Delete: orclindexedattribute
Orclindexedattribute: foo
```

**See Also:** "ldapmodify Syntax" on page A-13

### Indexing an Attribute for Which Directory Data Exists by Using the Catalog Management Tool

Use the Catalog Management Tool to index an attribute for which data already exists and to drop an index from an attribute.

**See:** "Catalog Management Tool Syntax" on page A-28

# 7

# Managing Directory Entries

This chapter explains how to view, add, and modify entries.

This chapter contains these topics:

- Managing Entries by Using Oracle Directory Manager
- Managing Entries by Using Command Line Tools
- Managing Entries by Using Bulk Tools
- Managing Entries with Attribute Options
- Managing Knowledge References (Referrals)

> **See Also:** Chapter 2, "Concepts and Architecture" for an overview of directory entries, directory information trees, distinguished names, and relative distinguished names

# Managing Entries by Using Oracle Directory Manager

This section contains these topics:



## Searching for Entries by Using Oracle Directory Manager

You can display all entries by using the navigator pane, or search for one or more specific entries by using the Oracle Directory Manager search feature.

To display an entry, in the navigator pane, expand Entry Management to display its subtree.

The root of the tree is listed first, then the second level, and so forth, moving from left to right. The subtree lists the **RDN** of each entry in hierarchical order. To see the lower level entries within any subtree, click the plus sign (+) to the left of the parent entry.

To search for a directory entry:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory_server_instance*, and select Entry Management. The Search fields appear in the right pane.
2. In the Root of the Search field, enter the **DN** of the root of your search.

   For example, suppose you want to search for an employee who works in the Manufacturing division in the IMC organization in the Americas. The DN of the root of your search would be:

   ```
   ou=Manufacturing,ou=Americas,o=IMC,c=US
   ```

   You would therefore type that DN in the Root of the Search text box.

You can also select the root of your search by browsing the **directory information tree (DIT)**. To do this:

a. Click Browse to the right of the Root of the Search field. The Select Distinguished Name (DN) Path: Tree View dialog box appears.

b. Click the plus sign (+) next to Tree View to display its entries.

c. Continue navigating to the entry that represents the level you want for the root of your search.

d. Select that entry, then click OK. The DN for the root of your search appears in the Root of the Search text box in the right pane.

3. In the Max Results (entries) box, type the maximum number of entries you want your search to retrieve. The default is 200.

4. In the Max Search Time (seconds) box, type the maximum number of seconds for the duration of your search. The value you enter here must be at least that of the default, namely, 25.

5. In the Search Depth list, select the level to which you want to search.

   The options are:

   - Base: Retrieves a particular directory entry. Along with this search depth, you use the Search criteria bar to select the attribute `objectClass` and the filter `Present`.
   - One Level: Limits your search to all entries beginning one level down from the root of your search
   - Subtree: Searches entries within the entire subtree, including the root of your search

6. In the Search Criteria box, use the lists and text fields on the search criteria bar to focus your search.

   a. From the list at the left end of the search criteria bar, select an attribute of the entry for which you want to search.

---

**Note:** Not all attributes are used in every entry. Be sure that the attribute you specify actually corresponds to one in the entry for which you are looking. Otherwise, the search will fail.

---

**b.** In the text box at the right end of the search criteria bar, type the value for the attribute you just selected. For example, if the attribute you selected was `cn`, you could type the particular common name you want to find.

**c.** From the list in the middle of the search criteria bar, select a filter. Options are:

| Filter | Description |
|---|---|
| Begins With | Searches by using only the first few characters of the attribute's value. For example, `cn Begins With Fran` retrieves all entries in which the first few letters of the `cn` attribute are `Fran`. These would include Frank, Fran, Frances, Franklin, etc. |
| Ends With | Searches for an entry by using only the last few characters of the specified attribute's value. For example, `cn Ends With son` retrieves Baldisson, Jacobson, Johnson, etc. |
| Contains | Searches for an entry in which the attribute you specified includes, but is not necessarily limited to, the value you enter. For example, `cn Contains Wins` retrieves all entries in which the `cn` attribute contains the letters `wins`. These would include Winslow, Czerwinski, Winship, etc. |
| Exact Match | Searches for an entry whose specified attribute is the same as the value you enter. For example, `cn Exactly Matches Franklin Baldwins` retrieves all entries in which the `cn` attribute has the value `Franklin Baldwins`. |
| Greater or Equal | Searches for an entry in which the specified attribute is numerically or alphabetically greater than or equal to the value you enter. For example, `cn Greater or Equal Frank` retrieves all entries with `cn` attributes that range from the first Frank to the end of the alphabet. |
| Less or Equal | Searches for entries in which the specified attribute is numerically or alphabetically less than or equal to the value you enter. For example, `cn Less or Equal Frank` retrieves all `cn` attributes from the first Frank to the beginning of the alphabet. |
| Present | Determines if an entry with the specified attribute is present at that level of the tree. You do not need to enter a value to use this relationship. The phrase `cn Present` retrieves all entries with the `cn` attribute at that level of the tree. |

7. To further refine your search, use the buttons in the Search Criteria box to enhance the search criteria bar.

| Button | Description |
|---|---|
| New | Creates a new search criteria bar in the Search Criteria field. This button is enabled only when the Search Criteria field is empty. |
| And | Creates another search criteria bar in the Search Criteria field. Matches all entries with one specified attribute with those that also have another specified attribute. For example, `cn=Baldwins And title=Laborer` retrieves all Baldwins who are also laborers. |
| Or | Creates another search criteria bar in the Search Criteria field. Matches all entries with either one specified attribute or another. For example, `title=Laborer Or title=Foreman` retrieves all employees who are either laborers or foremen. |
| Not | Negates the criterion in the selected search criteria bar and retrieves all entries that do not have the specified criterion. For example, `cn=Frank And Not title=Laborer` retrieves all persons named Frank who are not laborers. |
| Delete | Deletes a selected search criteria bar |

8. Click Search. The results of your search appear in the Distinguished Name box.

**See Also:** "Configuring Searches" on page 5-18 for instructions on setting the number of entries to display in searches, and to set the time limit for searches

## Searching for Audit Log Entries by Using Oracle Directory Manager

You can also search for audit log entries by using Oracle Directory Manager.

To use Oracle Directory Manager to view audit log entries:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory_server_instance*, and select Audit Log Management. The corresponding right pane appears.
2. Follow the instructions in "Searching for Entries by Using Oracle Directory Manager" on page 7-2 to search for particular types of entries in the audit log. The results of the search appear in the lower box.
3. To view the properties of a particular audit log entry, select it in the lower box, then click View Properties. The Audit Log Entry dialog box displays the properties for the audit log entry you selected.

**See Also:** "Configuring Searches" on page 5-18 for instructions on setting the number of entries to display in searches, and to set the time limit for searches

## Viewing Attributes by Using Oracle Directory Manager

Once you have displayed the results of your search, click the entry whose attributes you want to view. An Entry dialog box displays the attributes for that entry.

Some attributes can also be DNs. For example, one attribute for a given employee might be that employee's manager who, in turn, has a DN. In this case, when you display the Entry dialog box for the employee, you would see a Browse button next to the Manager text box. To find information about that manager, click Browse to display the Directory: Entry Management dialog box, then follow the steps mentioned in "Searching for Entries by Using Oracle Directory Manager" on page 7-2.

## Adding Entries by Using Oracle Directory Manager

**Note:** This release of Oracle Internet Directory does not support the adding of JPEG images by using Oracle Directory Manager. You may add a JPEG image by using the ldapadd command. For more information, see "Example: Adding a User Entry by Using ldapadd" on page 7-13.

### Adding a New Entry by Using Oracle Directory Manager

To add or delete entries with Oracle Directory Manager, you must have write access to the parent entry and you must know the DN for the new entry.

To add a new entry:

1. Expand Oracle Internet Directory Servers > *directory_server_instance*, then select Entry Management.
2. On the toolbar, click Create. The New Entry dialog box appears.
3. In the Distinguished Name field, type the full DN. You may also click Browse to locate and select the DN of the parent for the entry you want to add. The entry you select appears in the Distinguished Name field. To the left of that parent DN, type the RDN for your new entry, followed by a comma.
4. To specify the **object classes** for the new entry, next to the Object Classes box, click Add. The Super Class Selector dialog box appears.
5. In the Super Class Selector dialog box, select an object class, then click Select. As you select from the object class list, mandatory and optional attributes populate the windows in the tab pages in the lower half of the New Entry dialog box. You must enter values into the mandatory attributes fields. You are not required to enter values into the optional attributes fields.
6. When you have selected the object classes and provided values for the appropriate attributes, click OK.

### Adding an Entry by Copying an Existing Entry in Oracle Directory Manager

You can use Oracle Directory Manager to create a new entry by copying from an existing entry and changing its DN. When you do this, you should also change the attributes, such as name and address, so that they correspond to the new DN. To add an entry, you must have write access to its parent.

> **Tip:** You can find a template for the new DN by looking up other similar entries in the search pane.

To add an entry by copying an existing entry:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory_server_instance*, then select Entry Management. The Search pane appears. Use it to search for an entry that you want to use as a template.
2. Double-click an entry from those retrieved. The Entry dialog box for that entry appears. This entry will serve as your template in the Create Like pane.
3. In the Entry dialog box, click Create Like. A New Entry: Create Like dialog box appears.

4. Change critical fields to tailor this entry to the one that you want to create. You must always change the DN and the common name in this operation, or the pane will not save your new entry data. For example, if you create an entry for Henri Latrobe using the entry for Henri Latour as the template, then you have to change `cn=Henri Latour` in the DN to `cn=Henri Latrobe`. You also have to change the Henri Latour value in the common name attribute to Henri Latrobe, and any other attributes that must be unique, such as employee number and telephone number.
5. Click OK to save your changes.

> **See Also:** The online help for this dialog box for details about adding information into fields

### Example: Adding a User Entry by Using Oracle Directory Manager

In this example, we create a user named Anne Smith and assign her a password.

1. Login as the administrator.
2. Expand Oracle Internet Directory Services > directory_*server_instance*, and select Entry Management.
3. On the toolbar, click the Create button. The New Entry dialog box appears.
4. In the Distinguished Name field, type the full DN. You may also click the Browse button to locate the DN of the parent for this entry, then type the RDN, namely, `cn=Anne Smith`, followed by a comma, to the left of that parent DN.
5. To the right of the Object Classes box, Click Add. The Super Class Selector dialog box appears.
6. In the Super Class Selector dialog box, select the `person` object class, then click Select. This returns you to the New Entry dialog box.
7. In the New Entry dialog box, click the Optional Properties tab, and scroll to the `userPassword` window.
8. Type the password for Anne Smith.

### Adding Group Entries by Using Oracle Directory Manager

A group entry is one that contains a list of entries, for example, an e-mail list. You associate it with either the `groupOfNames` or `groupOfUniqueNames` object class, which has the object class `orclPrivilegeGroup` as a subclass.

You determine membership in the group by adding DNs to the multivalued attribute `member` if the entry belongs to the `groupOfNames` object class, or `uniqueMember` if the entry belongs to the `groupOfUniqueNames` object class.

To add a group entry:

1. Expand Oracle Internet Directory Servers > *directory_server_instance*, then select Entry Management.
2. On the toolbar, click Create. The New Entry dialog box appears.
3. In the Distinguished Name field, type the full DN. You may also use the Browse button to locate the DN of the parent for the entry you want to add, then type the RDN for the new entry, followed by a comma, to the left of that parent DN.
4. To specify the object classes you want to use for the new entry, to the right of the Object Classes box, click Add. The Super Class Selector dialog box appears.
5. In the Super Class Selector dialog box, select the `top` object class, then click the Select button. The `top` object class appears in the Object Classes box of the New Entry dialog box.
6. In the same way:
   a. To the right of the Object Classes box, click Add.
   b. From the Super Class Selector dialog box, select the `groupOfNames` or `groupOfUniqueNames` object class.
   c. Click Select. The object class you selected appears in the Object Classes window of the New Entry dialog box.
7. Enter the mandatory and optional attributes for your group entry.

   If you selected the `groupOfNames` object class, a Browse button appears next to some of the fields, for example, the member field on the Mandatory Properties tab page. To enter a mandatory property by browsing:

   a. Click Browse. The Directory: Entry Management dialog box appears.
   b. Use this dialog box to search for a particular entry you want to add to the list.

c. In the Distinguished Name window of the Directory: Entry Management dialog box, select the entry, then click OK. This returns you to the New Entry dialog box. The entry you just selected is added to the list in the members window.

8. Click Ok.

**See Also:**

- "Searching for Entries by Using Oracle Directory Manager" on page 7-2 for instructions on using the search pane
- "Privilege Groups" on page 9-4 for instructions on setting access control policies for group entries
- Access Control and Authorization on page 2-16 and Chapter 9, "Managing Directory Access Control" for information about access privileges

## Modifying Entries by Using Oracle Directory Manager

Oracle Directory Manager is governed by standard LDAP conventions, including the following:

- You cannot change object classes that are used by an entry once you have assigned object classes to that entry and populated its attributes with data.

  For example, if you configure an entry to use object classes `Person` and `Organizational Role`, you cannot later add another object class to this entry.

- You cannot add mandatory attributes to an object class already in use by some entries. You may add optional attributes to object classes that are already in use by entries. If you add optional attributes to an object class already in use by some entries, no special rules apply—they are added as empty attributes to those entries.

To modify an entry:

1. Perform a search for the entry you want to modify as described in "Searching for Entries by Using Oracle Directory Manager" on page 7-2.
2. In the Distinguished Name box of the right pane, select the entry you want to modify.
3. Click Edit. The Entry dialog box appears.

4. Select the Properties tab page. If you do not see the attributes you want to add or modify, then, at the top of the tab page, select View Properties: All.
5. In the Properties tab page, modify the values of any editable attributes.
6. Click OK.

### Example: Modifying a User Entry by Using Oracle Directory Manager

In this example, we modify the password for the entry we created for Anne Smith in the section "Example: Adding a User Entry by Using Oracle Directory Manager" on page 7-8.

1. Perform a search for the Anne Smith entry.
2. In the Distinguished Name box of the right pane, select the entry for Anne Smith.
3. Click Edit.
4. In the Entry dialog box, scroll to the `userPassword` window and modify the value.
5. Click OK.

## Managing Entries by Using Command Line Tools

This section points you to the command line tools you can use in managing entries. It also provides several examples of entry management by using command line tools. It contains these topics:

- Command Line Tools for Managing Entries
- Example: Adding a User Entry by Using ldapadd
- Example: Adding an Attribute Option
- Example: Modifying a User Entry by Using ldapmodify

## Command Line Tools for Managing Entries

The following table lists each of the command line tools, and tells you where to find syntax and usage notes for each one.

| Tool | Task(s) | Syntax and Usage Notes |
|---|---|---|
| ldapsearch | Search for directory entries. | "ldapsearch Syntax" on page A-18 |
| ldapbind | Authenticate a user or client to a directory server.<br>Verify that you can connect a client to a server. | "ldapbind Syntax" on page A-8 |
| ldapadd | Add entries one at a time.<br>Add new configuration set entries.<br>Configure a server with an input file. | "ldapadd Syntax" on page A-4 |
| ldapaddmt | Add several entries concurrently by using this multithreaded tool. | "ldapaddmt Syntax" on page A-6 |
| ldapmodify | Create, update, and delete attribute data for an entry.<br>Modify configuration set entries.<br>Modify DN or RDN of an entry. | "ldapmodify Syntax" on page A-13 |
| ldapmodifymt | Modify several entries concurrently by using this multithreaded tool. | "ldapmodifymt Syntax" on page A-16 |
| ldapdelete | Delete entries. | "ldapdelete Syntax" on page A-10 |
| ldapcompare | Compare attribute values you specify with those in a directory entry. | "ldapcompare Syntax" on page A-9 |
| ldapmoddn | Modify the DN or RDN of an entry.<br>Rename an entry or a subtree.<br>Move an entry or a subtree under a new parent. | "ldapmoddn Syntax" on page A-11 |

## Example: Adding a User Entry by Using ldapadd

The following example shows an LDIF file, named `entry.ldif`, for the user entry for an employee named John:

```
dn: cn=john, c=us
objectclass: top
objectclass: person
objectclass: organizationalPerson
objectclass: inetOrgPerson
cn: john
cn;lang-fr:Jean
cn;lang-en-us:John
sn: Doe
jpegPhoto: /photo/john.jpg
userpassword: welcome
```

This file contains the `cn`, `sn`, `jpegPhoto`, and `userpassword` attributes.

For the `cn` attribute, it specifies two options: `cn;lang-fr`, and `cn;lang-en-us`. These options return the common name in either French or American English.

For the `jpegPhoto` attribute, it specifies the path and file name of the corresponding JPEG image you want to include as an entry attribute.

## Example: Modifying a User Entry by Using ldapmodify

The following example changes the password for a user named Audrey from `welcome` to `audreyspassword`. As in the example above, the data for this user entry is in the `entry.ldif` file. This file contains the following:

```
dn: cn=audrey,c=us
changetype: modify
replace: userpassword
userpassword: audreyspassword
```

Issue this command to modify the file:

```
ldapmodify -p 389 -b  -f entry.ldif
```

# Managing Entries by Using Bulk Tools

This section lists and describes some of the more common tasks you perform with bulk tools.

This section contains these topics:

- Importing an LDIF File by Using bulkload
- Converting Directory Data to LDIF
- Modifying a Large Number of Entries
- Deleting a Large Number of Entries

> **See Also:** "Using Bulk Tools" on page 4-13 for an overview of these tools

## Importing an LDIF File by Using bulkload

To import an LDIF file, you use the bulkload utility. This section discusses the tasks to process an LDIF file through bulkload.

> **Note:** Before performing a bulk load, stop the Oracle Internet Directory processes. See Chapter 3, "Preliminary Tasks" for instructions on stopping directory server instances.

This section contains these topics:

- Task 1: Back Up the Oracle Server
- Task 2: Find Out the Oracle Internet Directory Password
- Task 3: Check Input for Schema and Data Consistency Violations
- Task 4: Generate the Input Files for SQL*Loader
- Task 5: Load the Input Files
- If Bulk Loading Fails

### Task 1: Back Up the Oracle Server

Before you import the file, back up the Oracle database server as a safety precaution.

> **See Also:** *Oracle8i Backup and Recovery Guide*

### Task 2: Find Out the Oracle Internet Directory Password

To use bulkload and the other shell script tools that have commands that end with `.sh`, you must provide the Oracle Internet Directory password. The default password is `ods`, although the system administrator can change it by using the **OID Database Password Utility**.

> **See Also:** "Using the OID Database Password Utility" on page 4-14

### Task 3: Check Input for Schema and Data Consistency Violations

On Solaris, the `bulkload.sh` file usually resides in `$`*`ORACLE_HOME`*`/ldap/bin`. On Windows NT, this file usually resides in *`ORACLE_HOME`*`\ldap\bin`.

Check the input file by typing:

```
bulkload.sh -connect net_service_name -check path_to_ldif-filename
```

All schema violations are reported in `$`*`ORACLE_HOME`*`/ldap/log/schemacheck.log`

If any violations are detected in the input file, use an ASCII text file editor to fix or remove them. If there are any duplicate entries, their DNs are logged in `$`*`ORACLE_HOME`*`/ldap/log/duplicate.log`.

### Task 4: Generate the Input Files for SQL*Loader

After you have fixed any errors in the input file, rerun bulkload with the `-generate` option as shown in the following example. During this step, LDIF data is converted to SQL*Loader specific format.

```
bulkload.sh -connect net_service_name -generate ldif-filename
```

All loading errors are reported in `$`*`ORACLE_HOME`*`/ldap/log`

When this command completes successfully, it generates `*.dat` files in the `$ORACLE_HOME/ldap/load` directory to be used by SQL*Loader in `-load` mode. Do not modify these files.

### Task 5: Load the Input Files

After you have generated the input files, rerun bulkload with the `-load` option. During this step, the `*.dat` files, which are in Oracle SQL*Loader specific format, are loaded into the database and the attribute indexes are created. The syntax is:

```
bulkload.sh -connect net_service_name -load
```

### If Bulk Loading Fails

All loading errors are reported in the $*ORACLE_HOME*`/ldap/log/directory` with the file extension `.bad`.
If bulk loading fails, the database could be left in an inconsistent state. It may be necessary to restore the database to its state prior to the bulk loading operation.

## Converting Directory Data to LDIF

Converting directory data to LDIF by using LDIF Writer makes the data available for loading into a new node in a replicated directory or into another node for backup storage.

**See Also:** "ldifwrite Syntax" on page A-27

## Modifying a Large Number of Entries

The bulkmodify utility enables you to modify a large number of existing entries in an efficient way.

**See Also:** "bulkmodify Syntax" on page A-25

## Deleting a Large Number of Entries

The bulkdelete utility enables you to delete an entire subtree efficiently.

**See Also:** "bulkdelete Syntax" on page A-22

# Managing Entries with Attribute Options

To manage entries with attribute options, you use command line tools. This section contains these topics:



## Example: Adding an Attribute Option

Suppose that you were adding the Spanish equivalent of an entry for John. As in the example above, the data for this user entry is in the `entry.ldif` file. This file contains the following:

```
dn: cn=john,c=us
changeType: modify
add: cn;lang-sp
cn;lang-sp: Juan
```

Issue this command to modify the file:

```
ldapmodify -p 389 -b  -f entry.ldif
```

## Example: Deleting an Attribute Option

The following example deletes the `cn;lang-fr` attribute option from the entry for John. As in the previous example, the data for this user entry is in the `entry.ldif` file. This file contains the following:

```
dn: cn=john, c=us
changetype: modify
delete: cn;lang-fr
cn;lang-fr: Jean
```

Issue this command to modify the file:

```
ldapmodify -p 389 -b  -f entry.ldif
```

### Example: Searching for Entries with Attribute Options

The following example retrieves entries with common name (`cn`) attributes that have an option specifying a language code attribute option. This particular example retrieves entries in which the common names are in French and begin with the letter R.

```
ldapsearch -p 389 -h myhost -b "c=US" -s sub "cn;lang-fr=R*"
```

Suppose that, in the entry for John, no value is set for the `cn;lang-it` language code attribute option. In this case, the following example fails:

```
ldapsearch -p 389 -h myhost -b "c=us" -s sub "cn;lang-it=Giovanni
```

**See Also:** "Attribute Options" on page 2-7

## Managing Knowledge References (Referrals)

A **knowledge reference**, also called a **referral**, is represented in the directory as a particular type of **entry**. When you create a knowledge reference entry, you associate it with the `referral` and `extensibleObject` **object classes**. Typically, you create knowledge reference entries at the place in the **DIT** where you want to establish the partition.

Knowledge references provide users with LDAP URLs. You enter these URLs as values for the `ref` attribute. There can be multiple `ref` attributes specified for any knowledge reference entry. Similarly, there can be multiple knowledge reference entries in the DIT.

**See Also:** "Distributed Directories: Partitioning" on page 2-42 for an overview of knowledge references and a description of **smart knowledge references** and **default knowledge references**

This section contains these topics:

- Configuring Smart Knowledge References
- Configuring Default Knowledge References

## Configuring Smart Knowledge References

When a user performs a search operation, Oracle Internet Directory looks for the knowledge reference entry within the specified scope of the search. If it finds the knowledge reference, then Oracle Internet Directory returns it to the client.

If a user performs an add, delete, or modify operation on an entry located below the knowledge reference entry, then Oracle Internet Directory returns the knowledge reference.

> **Note:** A search result can contain regular entries along with knowledge references.

For example, suppose you want to partition the DIT based on the geographical location of the directory servers. In this example, assume that:

- The `c=us` naming context is held locally on Server A and Server B in the United States.
- The `c=uk` naming context is held locally on Server C and Server D in the United Kingdom.

You would configure knowledge references between these two naming contexts as follows:

1. On Server A in the United States, configure a knowledge reference for the `c=uk` object on Server C and Server D:

```
dn: c=uk
c: uk
ref: ldap://host C:389/c=uk
ref: ldap://host D:686/c=uk
objectclass: top
objectclass: referral
objectClass: extensibleObject
```

2. Configure a similar knowledge reference on Server C in the United Kingdom for the `c=us` object on Server A and Server B:

```
dn: c=us
c: us
ref: ldap://host A:4000/c=us
ref: ldap://host B:5000/c=us
objectclass: top
objectclass: referral
objectClass: extensibleObject
```

Results:

- A client querying Server A with base `o=foo,c=uk` receives a knowledge reference
- A client querying Server C with base `o=foo,c=us` receives a knowledge reference
- An add operation of `o=foo,c=uk` on either Server A or Server B fails. Instead, Oracle Internet Directory returns a knowledge reference.

## Configuring Default Knowledge References

Oracle Internet Directory uses the `namingcontext` attribute in the **DSE** to determine all the **naming contexts** held locally by the server. Be sure that the `namingContext` attribute correctly reflects the naming context information.

You specify default knowledge references by entering a value for the `ref` attribute in the DSE entry. If the `ref` attribute is not in the DSE entry, then no default knowledge reference is returned.

When configuring a default knowledge reference, do not specify the DN in the LDAP URL.

For example, suppose that the DSE entry on Server A contains the following `namingContext` value:

```
namingcontext: c=us
```

Further, suppose that the default knowledge reference is:

```
Ref: ldap://host PQR:389
```

Now, suppose that a user enters an operation on Server A that has a base DN in the naming context `c=canada`, for example:

```
ou=marketing,o=foo,c=canada
```

This user would receive a knowledge reference to the host PQR. This is because Server A does not hold the `c=canada` base DN, and the `namingcontext` attribute in its DSE does not hold the value `c=canada`.

> **See Also:** "About Knowledge References (Referrals)" on page 2-43 for a conceptual discussion of knowledge references

# 8

# Managing Secure Sockets Layer (SSL)

This chapter explains how to configure the features of Secure Sockets Layer (SSL). If you use Secure Sockets Layer (SSL), you may also configure strong authentication, data integrity, and data privacy.

This chapter contains these topics:

- Supported Cipher Suites
- SSL Client Scenarios
- Configuring SSL Parameters
- Issues Specific to This Release of Oracle Internet Directory

> **See Also:** "Security" on page 2-12 for a conceptual overview of SSL in relation to Oracle Internet Directory

## Supported Cipher Suites

A cipher suite is a set of authentication, encryption, and data integrity algorithms used for exchanging messages between network nodes. During an SSL handshake, the two nodes negotiate to see which cipher suite they will use when transmitting messages back and forth.

The Oracle Internet Directory supports the following SSL cipher suites:

***Table 8–1 SSL Cipher Suites Supported in Oracle Internet Directory***

| Cipher Suite | Authentication | Encryption | Data Integrity |
|---|---|---|---|
| SSL_RSA_EXPORT_WITH_DES40_CBC_SHA | RSA | DES40 | SHA |
| SSL_RSA_EXPORT_WITH_RC4_40_MD5 | RSA | RC4_40 | MD5 |
| SSL_RSA_WITH_NULL_SHA | RSA | None | SHA |
| SSL_RSA_WITH_NULL_MD5 | RSA | None | MD5 |

## SSL Client Scenarios

Oracle Internet Directory clients can use SSL 2.0 or SSL 3.0. A client over SSL can connect to a server anonymously or by using either simple or strong authentication.

When both a client and server authenticate themselves to each other, SSL derives the identity information it requires from the X509v3 digital certificates.

## Configuring SSL Parameters

During start-up of a directory **server instance**, the directory reads a set of configuration parameters, including the parameters for the SSL profile. If you are going to run the directory with SSL enabled, you need to examine—and possibly reconfigure—the SSL parameters in the **configuration set entry**.

To run a server instance in secure mode, modify the configuration settings to run with the secure port 636 as the default port.

You can create and modify multiple sets of configuration parameters with differing values, using a different configuration set entry for each instance of Oracle Internet Directory. This is a useful way to accommodate clients with different security needs.

Oracle Corporation recommends that you create separate configuration sets and modify their SSL values, rather than modify SSL values in the default configuration

set. This is because the default configuration set may be required by Oracle Support Services in the diagnosis of certain technical issues.

**See Also:**

- "Managing Server Configuration Set Entries" on page 5-2 for instructions on how to set these parameters
- "Configuration Set Entry Attributes" on page E-5 for a description of these parameters

## Configuring SSL Parameters by Using Oracle Directory Manager

You can examine and modify the values for the SSL configuration parameters in each configuration set entry that you have created and in each server instance that is currently running.

---

**Note:** You cannot directly change the parameters for an active instance. If you want to change the parameters for an active instance, change the parameters in a configuration set entry and save it. After it is saved, you can stop current instances and refer to the newly modified configuration set in the start server message.

---

To view and modify SSL configuration parameters:

1. In Oracle Directory Manager's navigator pane, expand Oracle Internet Directory Servers > *directory server* > Server Management.
1. Expand either Directory Server or Replication Server, as appropriate. The numbered configuration sets are listed beneath your selection.
2. Select the configuration set that you want to examine. The group of tab pages for that configuration set entry appear in the right pane.
3. Select the SSL Settings tab page.

You can change the parameters in this tab page and save them. The fields in this tab page are described in the following table:

| Field | Description |
|---|---|
| SSL Enable | Select to enable SSL authentication. If you do not select this check box, SSL is not enabled, and you do not need to set any other parameters on this page. |
| SSL Authentication | Choose one of the following:<br>■ No SSL Authentication—Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, SSL encryption / decryption only is used.<br>■ SSL Client and Server Authentication—Both client and server authenticate themselves to each other and send certificates to each other.<br>■ SSL Server Authentication—Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |
| SSL Wallet URL | Type the location of the SSL wallet. If you elect to change the location of the Oracle wallet, you must change this parameter. You must set the wallet location on both the client and the server. For example, on Solaris, you could set this parameter as follows:<br>`orclsslwalleturl=file:/Home/my_dir/my_wallet`<br>On Windows NT, you could set this parameter as follows:<br>`file:C:\my_dir\my_wallet` |
| SSL Wallet Password | Type the password for the server-side wallet. This password was set during creation of the wallet. If you change the password, you must change this parameter. |
| SSL Wallet Confirm Password | Retype the new password in this field when you change the password. |
| SSL Port | The default SSL port is 636. You can change the SSL port. |

**See Also:** "Managing Server Configuration Set Entries by Using Oracle Directory Manager" on page 5-4 for information about changing parameters in a configuration set entry

### Configuring SSL Parameters by Using Command Line Tools

**See Also:** "Managing Server Configuration Set Entries by Using Command Line Tools" on page 5-10

## Issues Specific to This Release of Oracle Internet Directory

Oracle Internet Directory release 2.1.1, the Oracle directory replication server cannot communicate directly with SSL-enabled Oracle directory server instances.

If you intend to support both SSL and non-SSL clients on the same host, you need to configure two distinct server instances.

**See Also:** Chapter 5, "Managing an Oracle Directory Server" for instructions on how to configure server instances

# 9

# Managing Directory Access Control

This chapter provides an overview of access control policies and describes how to administer directory access control by using either Oracle Directory Manager or the command-line tool ldapmodify.

This chapter contains these topics:

- Overview of Access Control Policy Administration
- Managing Access Control by Using Oracle Directory Manager
- Managing Access Control by Using Command Line Tools

**See Also:**

- "Access Control and Authorization" on page 2-16 for a conceptual explanation before you begin implementing and administering access control policies
- Appendix D, "Using Access Control Directive Format" for information about the format or syntax of Access Control Items (ACIs)

# Overview of Access Control Policy Administration

You manage access control policies by configuring the values of the **ACI** attributes within appropriate entries. You can do this by using either Oracle Directory Manager or ldapmodify.

This section contains these topics:



## Access Control Management Constructs

This section contains these topics:



### orclACI

The `orclACI` attribute contains **Access Control List (ACL)** directives that are prescriptive in nature, that is, these directives apply to all entries in the subtree below the ACP where this attribute is defined. Any entry in the directory can contain values for this attribute. Access to this attribute itself is controlled in the same way as access to any other attribute.

**Note:** It is possible to represent ACL directives specific to a single entry in the `orclACI` attribute. However, in such scenarios, for administrative convenience and performance advantages, Oracle Corporation recommends using `orclEntryLevelACI`—discussed in "orclEntryLevelACI" on page 9-3. This is because the LDAP operational overhead increases with the number of directives represented through `orclACI`. You can reduce this overhead by moving entry specific directives from `orclACI` to `orclEntryLevelACI`.

### Access Control Policy Points (ACPs)

ACPs are entries in which the `orclACI` attribute has been given a value. The `orclACI` attribute value represents the access policies that are inherited by the subtree of entries starting with the ACP as the root of the subtree.

When a hierarchy of multiple ACPs exists in a directory subtree, the subordinate entries in that subtree inherit the access policies from all of the ACPs that are superior to the entry. The resulting policy is an aggregation of the policies within the ACP hierarchy above the entry.

For example, if an ACP is established in the HR department entry, and the Benefits, Payroll, and Insurance groups are entries within the HR department, then any entry within those groups inherits the access rights specified in the HR department entry.

When there are conflicting policies within a hierarchy of ACPs, the directory applies well-defined precedence rules in evaluating the aggregate policy.

**See Also:** "How ACL Evaluation Works" on page 9-10

### orclEntryLevelACI

When a policy pertains only to a specific entity—for example, a special user—you can maintain, within a single entry, the ACL directives specific to that entry. Oracle Internet Directory enables you to do this through a user-modifiable operational attribute called `orclEntryLevelACI`. The `orclEntryLevelACI` attribute contains ACL directives that apply only to the entry with which it is associated.

Any directory entry can optionally carry a value for this attribute. This is because Oracle Internet Directory extends the abstract class `top` to include `orclEntryLevelACI` as an optional attribute.

The orclEntryLevelACI attribute is multi-valued and has a structure similar to that of orclACI. The structure definition is provided later in this chapter.

### Privilege Groups

Group entries in Oracle Internet Directory are associated with either the groupOfNames or the groupOfUniqueNames object class. Membership in the group is specified as a value of the member or uniqueMember attribute respectively.

It is possible to specify access rights for a group of people or entities. Such groups are called privilege groups and are associated with the orclPrivilegeGroup object class.

To grant access rights to a group of users, you create a group entry in the usual way, then associate it with the orclPrivilegeGroup object class. You then specify the access policies applicable to that group.

Entries can have either direct memberships to groups, or indirect memberships to other groups by means of nested groups, thus forming a forest of privilege groups. Access policies specified at a given level are applicable to all the members directly or indirectly below it.

Because Oracle Internet Directory evaluates for access control purposes only groups marked as privilege groups, it does not allow setting access policies for non-privilege groups. When a user binds with a specific distinguished name (DN), Oracle Internet Directory computes the user's direct membership in privilege groups. Once it knows the first level groups for the given DN, Oracle Internet Directory computes nesting of all these first level groups into other privilege groups. This process continues until there are no more nested groups to be evaluated.

It is imperative that all groups created for access control purposes, nested or otherwise, be marked as privilege groups by associating them with the orclPrivilegeGroup object class. A normal group will not be considered for access control purposes even though it may be a member of a privilege group.

For example, consider the following group of entries each of which, with the exception of group4, is marked as a privilege group (`objectclass:orclprivilegegroup`). One can set access control policies such that they are applicable to the members of group1, group2, and group3.

```
dn: cn=group1, c=us
cn: group1
objectclass: top
objectclass: groupofUniquenames
objectclass: orclprivilegegroup
uniquemember:  cn=john smith, c=us
uniquemember:  cn=joe smith, c=us
uniquemember:  cn=bill smith, c=us
```

```
dn: cn=group2, c=us
cn: group2
objectclass: top
objectclass: groupofUniquenames
objectclass: orclprivilegegroup
uniquemember:  cn=john smith, c=us
uniquemember:  cn=joe smith, c=us
uniquemember:  cn=bill smith, c=us
```

```
dn: cn=group3, c=us
cn: group3
objectclass: top
objectclass: groupofUniquenames
objectclass: orclprivilegegroup
uniquemember:  cn=group2, c=us
uniquemember:  cn=group1, c=us
uniquemember: cn=group4, c=us
```

```
dn: cn=group4, c=us
cn: group4
objectclass: top
objectclass: groupofUniquenames
uniquemember:  cn=john smith, c=uk
uniquemember:  cn=joe smith, c=uk
uniquemember cn=bill smith, c=us
```

Group `cn=group3,c=us` contains the following nested groups:

- `cn=group2,c=us`
- `cn=group1,c=us`
- `cn=group4,c=us`

Access control policies for group3 are applicable to members of group3, group1, and group2 because each of them is marked as a privilege group. These same access control policies are not applicable to the members of group4 because group4 is not marked as a privilege group.

For example, suppose that the user binds to Oracle Internet Directory as a member of group 4 with the DN `cn=john smith,c=uk`. None of the access policies applicable to the members of group3 will apply to this user. This is because his only direct membership is to a non-privilege group. By contrast, if the user were to bind as `cn=john smith,c=us`—that is, as a member of group1 and group2—then his access rights will be governed by access policies set up for members of group1, group2, as well as group3 (in which group1 and group2 are nested). This is because all three groups are associated with the object class `orclPrivilegeGroup`.

## Access Control Information Components

Access Control Information associated with a directory object represents the permissions on the given object that various directory user entities (or subjects) have. Thus, an ACI consists of three components:

- Object: To What Are You Granting Access?
- Subject: To Whom Are You Granting Access?
- Operations: What Access Are You Granting?

### Object: To What Are You Granting Access?

The *object* part of the access control directive determines the entries and attributes to which the access control applies. It can be either an entry or an attribute. Entry objects associated with an ACI are implicitly identified by the entry or the subtree where the ACI itself is defined. Any further qualification of objects at the level of attributes is specified explicitly in the ACL expressions.

In the `orclACI` attribute, the entry DN component of the object of the ACI is implicitly that of all entries within the subtree starting with the ACP as its topmost entry. For example, if `dc=com` is an ACP, then the directory area governed by its ACI is:

```
.*, dc=com.
```

However, since the directory area is implicit, the DN component is neither required nor syntactically allowed.

In the `orclEntryLevelACI` attribute, the entry DN component of the object of the ACL is implicitly that of the entry itself. For example, if `dc=acme,dc=com` has an entry level ACI associated with it, the entry governed by its ACI is exactly: `dc=acme,dc=com`. Since it is implicit, the DN component is neither required nor syntactically allowed.

The object portion of the ACL allows entries to be optionally qualified by a filter matching some attribute(s) in the entry:

`filter=(`*`ldapFilter`*`)`

where *`ldapFilter`* is a string representation of an LDAP search filter. The special entry selector * is used to specify all entries.

Attributes within an entry are included in a policy by including a comma-separated list of attribute names in the object selector.

`attr=(`*`attribute_list`*`)`

Attributes within an entry are excluded from a policy by including a comma-separated list of attribute names in the object selector.

`attr!=(`*`attribute_list`*`)`

---

**Note:** Access to the entry itself must be granted or denied by using the special object keyword ENTRY. Note that giving access to an attribute is not enough; access to the entry itself through the ENTRY keyword is necessary.

---

**See Also:**

- "Examples: Managing Access Control" on page 9-35 for examples using command line tools
- Appendix D, "Using Access Control Directive Format" for information about the format or syntax of ACI

### Subject: To Whom Are You Granting Access?

This section describes the authentication mode, called the bind mode, used to verify the identity of the subject, also called the entity, to whom access is granted.

**Bind Mode** The bind mode specifies the method of authentication to be used by the subject. There are four modes:

- `Simple`: Simple password-based authentication
- `SSLNoauth`: For SSL-based clients with either anonymous or simple password based authentication

- SSLOneway: For SSL-based clients with server authentication with either anonymous or password based authentication
- SSLTwoway: For SSL-based clients with strong authentication through SSL.

The bind mode is optional in subject specification. When specified, the mode should match the mode specified in the ACI.

**Entity** The entity component identifies the entity or entities being granted access. Note that access is granted to entities, not entries.

Entities can be specified by:

- The special "*" identifier, matching any entry
- The keyword SELF matching the entry protected by the access
- A regular expression matching an entry's distinguished name: dn=*regex*
- The members of a privilege group object: group=*dn*
- An entry listed in a DN-valued attribute in the entry to which the access applies: dnattr=(*dn-valued_attribute_name*)

The dnattr specification is used to give access to a group entry to whomever is listed as the owner of the group entry.

### Operations: What Access Are You Granting?

The kind of access granted can be one of the following:

- None
- Compare / nocompare
- Search / nosearch
- Browse / nobrowse
- Read / noread
- Selfwrite / noselfwrite
- Write / nowrite
- Add / noadd
- Delete / nodelete

Note that each access level can be independently granted or denied. The no*xxx* means *xxx* permission is denied.

| Access Level | Description |
|---|---|
| None | No access rights. The effect of granting no access rights to a subject-object pair is to make the directory appear to the subject as though the object were not present in the directory. |
| Add | Right to add entries under a target directory entry |
| Browse | Permission to return the DNs in the search result. It is equivalent to the list permission in X.500. This permission is also required for a client to use an entry DN as the base DN in an ldapsearch operation. |
| Compare | Right to perform compare operation on the attribute value |
| Delete | Right to delete the target entry |
| Read | Right to read attribute values. Even if read permission is available for an attribute, it cannot be returned unless there is browse permission on the entry itself. |
| Search | Right to use an attribute in a search filter |
| Selfwrite | Right to add oneself to, delete oneself from, or modify one's own entry in a list of DNs group entry attribute. Use this to allow members to maintain themselves on lists. For example, the following command allows people within a group to add or remove only their own DN from the member attribute:<br>`access to attr=(member) by dnattr=(member) (selfwrite)`<br>The `dnattr` selector indicates that the access applies to entities listed in the member attribute. The `selfwrite` access selector indicates that such members can add or delete only their own DN from the attribute. |
| Write | Right to modify / add / delete the attributes of an entry. |

Note that some access permissions are associated with entries and others with attributes.

| Permissions for Entries: | Permissions for Attributes: |
|---|---|
| Browse / nobrowse | Compare / nocompare |
| Add / noadd | Search / nosearch |
| Delete / nodelete | Read / noread |
| None | Selfwrite / noselfwrite |
| | Write / nowrite |
| | None |

The entry level access directives are distinguished by the keyword ENTRY in the object component.

**Note:** By default, for both structural and content access items, everyone is given access to read, search, write, and compare all attributes in an entry, and selfwrite permissions are unspecified. If an entry is unspecified, access is determined at the next highest level in which access is specified.

## How ACL Evaluation Works

This section contains these topics:

- About ACL Evaluation
- ACL Evaluation Precedence Rules
- Assigning More Than One ACI to the Same Object
- Granting Exclusionary Access to Objects
- ACL Evaluation For Groups

## About ACL Evaluation

When processing a request, the access level granted to the requester has to be evaluated for each of the attributes involved in the request. This evaluation is done systematically for each attribute associated with every entry involved in an LDAP operation.

The process of evaluating access to any object (attribute in an entry) involves potentially examining all the ACI directives that are applicable for that object. This is because of the hierarchical nature of ACPs and the inheritance of policies from superior ACPs to subordinate ACPs.

The evaluation starts with examining ACI directives in the entry's entry level ACI, `orclEntryLevelACI`. Until the evaluation is complete, the ACP policies are successively considered, starting with the immediate ACP, followed by the chain of its superior ACPs.

The access evaluation is done for the entry and each of its attributes individually. Oracle Internet Directory evaluates entry level access permissions to see whether the given subject is allowed to perform the given operation.

During ACL evaluation, an attribute is said to be in one of the following states:

| State | Description |
|---|---|
| Resolved with permission | The required access for the attribute has been granted in the ACI. |
| Resolved with denial | The required access for the attribute has been explicitly denied in the ACI. |
| Unresolved | No applicable ACI has yet been encountered for the attribute in question. |

For all operations except search, the evaluation stops if:

- Access to the entry itself is denied
- Any of the attributes reach the resolved with denial state.

In this case the operation would fail and an error would be returned to the client.

For a search operation, the evaluation continues until all the attributes reach the resolved state. Attributes that are resolved with denial are not returned.

### ACL Evaluation Precedence Rules

An LDAP operation requires the BindDN, or subject, of the LDAP session to have certain permissions to the objects affected by the operation—including permissions on the entry itself and on the individual attributes of the entry.

Typically, there could be a hierarchy of access control administration authorities, starting from the root of a naming context down to successive administrative points (or access control policy points). An ACP is any entry which has a defined value for the `orclACI` attribute. Additionally, the access information specific to a single entry can also be represented within the entry itself (`orclEntryLevelACI`).

ACL evaluation involves determining whether a subject has sufficient permissions to perform an LDAP operation. Typically an `orclentryLevelACI` or `orclACI` might not contain all the necessary information for ACL evaluation. Hence, all available ACL information is processed in a certain order until the evaluation is fully resolved:

- The entry level ACI is examined first. ACI in the `orclACI` are examined starting with the ACP closest to the target entry and then its superior ACP and so on.
- At any point, if all the necessary permissions have been determined, the evaluation stops; otherwise, the evaluation continues.
- Within a single ACI, if the entity associated with the session DN matches more than one item identified in the *by* clause, the effective access evaluates to:
  - The union of all the granted permissions in the matching by clause items

    ANDed with
  - The union of all the denied permissions in the matching by clause items

**Precedence at the Entry Level** ACIs at the entry level are evaluated in the following order:

1. With a filter. For example:

```
access to entry filter=(cn=p*)
    by group1 (browse,add,delete)
```

2. Without a filter. For example:

```
access to entry
    by group1 (browse,add,delete)
```

**Precedence at the Attribute Level** At the attribute level, specified ACIs have precedence over unspecified ACIs.

*Specified* ACIs at the attribute level are evaluated in the following order:

1. Those with a filter. For example:

```
access to attr=(salary) filter=(salary >10000)
    by group1 (read)
```

2. Those without a filter. For example:

```
access to attr=(salary)
    by group1 (search,read)
```

*Unspecified* ACIs at the attribute level are evaluated in the following order:

1. With a filter. For example:

```
access to attr=(*) filter (cn=p*)
    by group1 (read,write)
```

2. Without a filter. For example:

```
access to attr=(*)
    by group1 (read,write)
```

### Assigning More Than One ACI to the Same Object

If there are two or more ACIs at the same ACP for the same object, then only one ACI is checked, and all other ACIs are ignored. For example, suppose you have the following two ACIs at the same ACP for the same entry:

- ACI #1:

```
access to entry
       by dn="cn=admin, dc=us,dc=acme,dc=com" (browse, add, delete)
```

- ACI #2:

```
access to entry
       by dn="cn=manager,dc=us,dc=acme,dc=com" (search, read)
```

If ACI #2 happens to be checked first, then the access granted specifically to the administrator in ACI #1 is ignored. If an administrator should then seek access to the entry, that access could not be resolved at this level of the hierarchy. The evaluation would have to move progressively up the hierarchy in search of resolution. If no resolution is found, all access is denied.

The solution is to create only one ACI at the same ACP for this entry. For example:

```
access to entry
    by dn="cn=admin, dc=us,dc=acme,dc=com" (browse, add, delete)
    by dn="cn=manager,dc=us,dc=acme,dc=com" (search, read)
```

Similarly, at the attribute level, suppose you have the following two ACIs:

- ACI #1:

  ```
  access to attr=(userpassword)
          by dnattr=(".*,dc=us,dc=acme,dc=com) (none)
  ```

- ACI #2:

  ```
  access to attr=(userpassword)
          by self (read,write)
  ```

If ACI #1 happens to be returned first, it wins, and the access granted to self in ACI #2 is ignored. If a user then wishes to change his or her own password, that access cannot be granted.

As with the ACIs for entries, the solution is to create only one ACI at the same ACP for this attribute. For example:

```
access to attr=(userpassword)
    by dnattr=(".*,dc=us,dc=acme,dc=com) (none)
    by self (read,write)
```

### Granting Exclusionary Access to Objects

If an ACI exists for a given object, and you want to specify access to all other objects except that one, then you must verify that the specified objects do not intersect. For example, suppose you have the following two ACIs:

- ACI #1:

  ```
  access to attr=(userpassword)by group1 (read,write)
  ```

- ACI #2:

  ```
  access to attr=(*)by group2 (read)
  ```

In this case, the two ACIs intersect, that is, both ACIs try to grant access to the `userpassword` attribute, but ACI #2 is unsuccessful. The reason is that, during the evaluation process, ACI #1 wins because, as noted in "ACL Evaluation Precedence Rules" on page 9-12, it is specified. This means that anyone in group2 who tries to access the `userpassword` attribute is not given access at this level of the hierarchy. The evaluation would have to move progressively up the hierarchy in search of resolution. If no resolution is found, all access is denied.

The solution is to use the following syntax for ACI #1 and ACI #2:

- ACI#1:

  ```
  access to attr=(userpassword)by group1 (read,write)by group2 (read)
  ```

- ACI #2:

  ```
  access to attr!=(userpassword)by group2 (read)
  ```

In the revised ACI #1, we give to group2 read access to the `userpassword` attribute.

In the revised ACI #2, we negate group2 access to the `userpassword` attribute, and we grant read access to all attributes *except* the `userpassword` attribute.

### ACL Evaluation For Groups

If an operation on an attribute or the entry itself is explicitly denied at an ACP low in the DIT, then, typically, the ACL evaluation for the attribute (or entry) is considered "Resolved with Denial." However, if the user of the session (bindDN) is a member of a group object, then the evaluation continues as if it is still unresolved. If permissions are granted to the user of the session at an ACP higher in the tree through a group subject selector, then such grants have higher precedence than any denials lower in the tree.

This scenario is the only case in which ACL policy at a higher level ACP has a higher precedence than that of an ACP lower in the DIT.

### Access Level Requirements for LDAP Operations

The following table lists LDAP operations and the access required to perform each one.

| Operation | Required Access |
|---|---|
| Create an object | Add access to the parent entry |
| Modify | Write access to the attributes that are being modified |
| ModifyDN | Delete access to the current parent and Add access to the new parent. |
| ModifyDN (RDN) | Write access to the naming attribute, that is, the RDN attribute |
| Remove an object | Delete access to the object being removed |
| Compare | Compare access to the attribute |
| Search | ■ Search access on the filter attributes and browse access on the entry (if only the entry DN needs to be returned as a result)<br>■ Search access on the filter attributes, browse access on the entry, and read permission on the attributes (for all attributes whose values need to be returned as a result) |

## Managing Access Control by Using Oracle Directory Manager

You can view and modify access control information configured within ACPs by using either Oracle Directory Manager or command line tools. This section explains how to accomplish these tasks by using Oracle Directory Manager.

**Note:** Immediately after installing Oracle Internet Directory, be sure to reset the default security configuration as described in "Task 3: Reset the Default Security Configuration" on page 3-9

This section contains these topics:

- Configuring the Display of ACPs in Oracle Directory Manager
- Configuring Searches for ACPs When Using Oracle Directory Manager
- Viewing an ACP by Using Oracle Directory Manager
- Modifying Existing ACPs and their ACI Directives by Using Oracle Directory Manager
- Adding an ACP and Creating Access Items by Using Oracle Directory Manager
- Example: Managing ACPs by Using Oracle Directory Manager
- Granting Entry-Level Access by Using Oracle Directory Manager

> **See Also:** Appendix A, "Syntax for LDIF and Command Line Tools" for a description of command line tools

## Configuring the Display of ACPs in Oracle Directory Manager

Oracle Directory Manager enables you to determine whether the navigator pane displays all ACPs automatically or only as the result of a search. If you have a large number of ACPs, you may want to display them only as the result of a search.

To configure the display of ACPs:

1. In the navigator pane, expand Oracle Internet Directory Servers and select the server you want to configure.
2. On the toolbar, click User Preferences. The User Preferences dialog box appears.
3. Select the Configure Access Control Policy Management tab page.
4. In the Configure Access Control Policy Management tab page, select either:
   - Always display all ACPs
   - Only display ACPs based on search request
5. Click Ok.

---

**Note:** To effect your changes, you must restart Oracle Directory Manager.

---

## Configuring Searches for ACPs When Using Oracle Directory Manager

For ACP searches, Oracle Directory Manager enables you to specify:

- The root of the search
- The maximum number of entries retrieved
- The time limit of the search
- The search depth.

To configure searches for ACP entries:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory_server*, and select Access Control Management.
2. On the toolbar, click Configure ACPs Search. The Configure ACPs Search dialog box appears.
3. In the Root of the Search field, enter the DN of the root of your search, or click Browse to navigate to it.
4. In the Max Results (entries) field, enter the number of entries you want ACP searches to retrieve.
5. In the Max Search Time (seconds) field, enter the maximum number of seconds for the duration of the search.
6. In the Search Depth list, select the level at which you want to search. Options are:
   - One Level—To limit the search to all ACP entries one level down from the root of the search
   - Subtree—To search entries within the entire subtree, including the root of the search
7. Click Ok.

## Viewing an ACP by Using Oracle Directory Manager

If you configured Oracle Directory Manager to display ACPs only as the result of a search, as described in "Configuring the Display of ACPs in Oracle Directory Manager" on page 9-17, then you can locate and view an ACP as follows:

1. Expand Oracle Internet Directory Servers > *directory_server*, then select Entry Management. Perform a search for the entry designated as an ACP. The search result appears in the Distinguished Name box in the lower half of the right pane.
2. In the Distinguished Name box, double-click the entry. The corresponding Entry dialog box appears.
3. To view subtree access controls for this ACP, select the Subtree Access tab.

   To view entry level access controls for this ACP, select the Local Access tab.

If you configured Oracle Directory Manager to always display ACPs, as described in "Configuring the Display of ACPs in Oracle Directory Manager" on page 9-17, then you can locate and view an ACP as follows:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory_server* > Access Control Management. All of the defined Access Control Policy Points (ACPs) appear both below Access Control Management in the navigator pane and in the right pane.
2. In the navigator pane, under Access Control Management, select an ACP to display its information in the right pane.

   You can alternatively double-click an ACP in the right pane to display the data in its own window.

The three fields in the Access Control Management pane are:

| Field | Description |
|---|---|
| Path to the Subtree Access Control Point | Contains the path defined by the ACP. If you have navigated down a tree to this point, the path to this point appears in this field. If you are creating a new ACP, you must enter the path to it here. |
| Structural Access Items (Entry Level Operations) | Lists access to entries. Items listed in the Structural Access Items box identify an entry by the following categories:<br>■ By Whom: To whom or what you are granting access (the subject)<br>■ Bind Mode: Whether bind mode (authentication) is used<br>■ Access rights: Browse, Add, and Delete<br>**See Also:** "Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-25 for instructions on how to modify structural access items |
| Content Access Items (Attribute Level Operations) | Lists items related to attributes within the entry or entries identified in the Entry Filter column. Columns in this window include:<br>■ By Whom: To whom or what you are granting access (the subject)<br>■ Bind Mode: Whether bind mode (authentication) is used<br>■ Op: The matching operation to be performed against the attribute. Choices are EQ (=) and NEQ (!=)<br>■ Attribute: The specific attribute to which access is granted or denied (the object)<br>■ Access rights: Read, Search, Write, Selfwrite, or Compare access<br>**See Also:** "Modifying Content Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-28 for instructions on how to modify content access items. |

## Modifying Existing ACPs and their ACI Directives by Using Oracle Directory Manager

ACPs are entries that contain prescriptive, that is, inheritable, access control information. This information affects the entry itself and all entries below it. You will most likely create ACPs to broadcast large-scale access control throughout a subtree.

This section contains these topics:

- Adding Structural Access Items to an Existing ACP by Using Oracle Directory Manager
- Adding Content Access Items to an Existing ACP by Using Oracle Directory Manager
- Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager
- Modifying Content Access Items of an Existing ACP by Using Oracle Directory Manager

### Adding Structural Access Items to an Existing ACP by Using Oracle Directory Manager

If you configured Oracle Directory Manager to always display ACPs, as described in "Configuring the Display of ACPs in Oracle Directory Manager" on page 9-17, then you can navigate to and add a structural access item to an existing ACP as follows:

1. In the navigator pane, expand Oracle Internet Directory > *directory server* > Access Control Management. Select Access Control Management. All of the defined Access Control Policy Points (ACPs) appear in a list below Access Control Management. They also appear in the right pane.
2. In the navigator pane, under Access Control Management, select an ACP to display its information in the right pane.
3. Just below the Structural Access Items box, click Create. The Structural Access Items dialog box displays three tabs: Entry Filter, By Whom, and Access Rights.
4. Use the Entry Filters tab page to narrow the set of entries to which you are granting access. If you want all entries below the ACP to be governed by the ACP, you do not need to use this tab page.

   You might choose an entry based on one or more attributes. For example, you might choose to search for all those whose title is administrative assistant, or for all those whose title is manager and whose organization unit is Americas.

In the Criteria box of the Entry Filters tab page, use the search criteria bar to select an attribute, enter a value for that attribute, and specify a filter for matching the specified attribute with the value you entered. To do this:

**a.** In the menu at the left end of the bar, select an attribute.

**b.** In the menu in the middle of the bar, select one of the following filter options:

| Filter | Description |
|---|---|
| Begins With | Searches by using only the first few characters of the attribute value |
| Ends With | Searches for an entry by using only the last few characters of the specified attribute value |
| Contains | Searches for an entry in which the attribute you specified includes, but is not necessarily limited to, the value you enter |
| Exact Match | Searches for an entry whose specified attribute is the same as the value you enter |
| Greater or Equal | Searches for an entry in which the specified attribute is numerically or alphabetically greater than or equal to the value you enter. An entry is alphabetically greater if it is closer to the beginning of the alphabet. |
| Less or Equal | Searches for entries in which the specified attribute is numerically or alphabetically less than or equal to the value you enter. An entry is alphabetically less if it is closer to the beginning of the alphabet. |
| Present | Determines if an entry with the specified attribute is present at that level of the tree. You do not need to enter a value to use this relationship. For example, the phrase `cn Present` retrieves all entries with a `cn` attribute value at that level of the tree. |

**c.** In the text field at the right end of the search criteria bar, type the value for the attribute you selected.

**5.** Select the By Whom tab page to define the subject of the ACI.

**a.** Specify the type of authentication—called Bind Mode—to be used by the subject (that is, the entity that seeks access). The bind mode is optional in subject specification. However, for the directive to be applicable, the bind mode specified on one node should match the bind mode specified on the node with which it is communicating.

There are five bind modes from which to select:

| Bind Mode | Description |
|---|---|
| None | No authentication |
| SSL No Authentication | Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, only SSL encryption / decryption is used. |
| SSL One Way | Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |
| SSL Two Way | Both client and server authenticate themselves to each other. They do this by sending certificates to each other. |
| Simple | The client identifies itself to the server by means of a DN and a password which are sent in the clear over the network. The server verifies that the DN and password sent by the client matches the DN and password stored in the directory. |

Specify the entity or entities to whom you are granting access. Options are:

| Entity | Description |
|---|---|
| Everyone (*) | All who try to access the entry |
| A Specific Group | A previously defined group name |
| A Specific Entry | A previously defined directory entry |
| A Subtree | An entire subtree in the directory, which you select |
| When Session User's Distinguished Name (DN) Is Identified by Attribute | Anyone whose DN is an attribute in the entry. For example, you might want to grant read access to a group entry to members of the group |
| When Session User's Distinguished Name (DN) Matches the Accessed Entry | Anyone who has correctly logged in as the entry specified |

**b.** Click OK.

**6.** Select the Access Rights tab page.

a. Select the appropriate options to specify the kinds of rights you want to grant: Browse, Add, or Delete.

b. Click Ok to close the Structural Access Items dialog box and return to the main Oracle Directory Manager window. The structural ACI you just set is listed in the Structural Access Items window of the main Oracle Directory Manager dialog box.

### Adding Content Access Items to an Existing ACP by Using Oracle Directory Manager

If you configured Oracle Directory Manager to always display ACPs, as described in "Configuring the Display of ACPs in Oracle Directory Manager" on page 9-17, then you can navigate to and add a content access item to an existing ACP as follows:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server* > Access Control Management. Select Access Control Management. All of the defined Access Control Policy Points (ACPs) appear in a list below Access Control Management in the navigator pane. They also appear in the right pane.
2. Select an ACP under Access Control Management in the navigator pane to display its information in the right pane, or double-click an ACP in the right pane to display the data in its own dialog box.
3. In the Content Access Items window, select the Content Access Item you want to modify.
4. Just below the Content Access Item box, click Create. The Content Access Items dialog box appears.
5. In the Entry Filter tab page, specify the items (if applicable) as described in "Adding Structural Access Items to an Existing ACP by Using Oracle Directory Manager" on page 9-21.
6. Select the By Whom tab page and specify the items as described in "Adding Structural Access Items to an Existing ACP by Using Oracle Directory Manager" on page 9-21.
7. Select the Attribute tab page.
   a. From the right list, select the attribute to which you want to grant or deny access.
   b. From the left list, select the matching operation to be performed against the attribute. Choices are EQ (Equal (=)) and NEQ (Not Equal (!=)).

8. Select the Access Rights tab page and specify the items as described in the section "Adding Structural Access Items to an Existing ACP by Using Oracle Directory Manager" on page 9-21.
9. Click OK.

### Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager

If you configured Oracle Directory Manager to always display ACPs, as described in "Configuring the Display of ACPs in Oracle Directory Manager" on page 9-17, then you can navigate to and modify a structural access item to an existing ACP as follows:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server* > Access Control Management. Select Access Control Management. All of the defined Access Control Policy Points (ACPs) appear in a list below Access Control Management in the navigator pane. They also appear in the right pane.
2. In the navigator pane, under Access Control Management, select an ACP to display its information in the right pane.
3. In the Structural Access Items window, select the item you want to modify, and, just below the Structural Access Items window, click Edit. The Structural Access Item dialog box appears.
4. Use the Entry Filters tab page to narrow the set of entries to which you are granting access. If you want all entries below the ACP to be governed by the ACP, proceed to the next step.

   You might choose an entry based on one or more attributes. For example, you might choose to search for all those whose title is secretary, or for all those whose title is manager and whose organization unit is Americas.

In the Criteria window of the Entry Filters tab page, use the search criteria bar to select an attribute, enter a value for that attribute, and specify a filter for matching the specified attribute with the value you entered. To do this:

**a.** In the menu at the left end of the bar, select an attribute.

**b.** In the menu in the middle of the bar, select one of the following filter options:

| Filter | Description |
|---|---|
| Begins With | Searches by using only the first few characters of the attribute value |
| Ends With | Searches for an entry by using only the last few characters of the specified attribute value |
| Contains | Searches for an entry in which the attribute you specified includes, but is not necessarily limited to, the value you enter |
| Exact Match | Searches for an entry whose specified attribute is the same as the value you enter |
| Greater or Equal | Searches for an entry in which the specified attribute is numerically or alphabetically greater than or equal to the value you enter. An entry is alphabetically greater if it is closer to the beginning of the alphabet. |
| Less or Equal | Searches for entries in which the specified attribute is numerically or alphabetically less than or equal to the value you enter. An entry is alphabetically less if it is closer to the beginning of the alphabet. |
| Present | Determines if an entry with the specified attribute is present at that level of the tree. You do not need to enter a value to use this relationship. For example, the phrase `cn Present` retrieves all entries with a `cn` attribute value at that level of the tree. |

**c.** In the text box at the right end of the search criteria bar, type the value for the attribute you selected.

5. Select the By Whom tab page.

   a. Specify the type of authentication—called Bind Mode—to be used by the subject (that is, the entity that seeks access). There are five bind modes from which to select:

| Bind Mode | Description |
|---|---|
| None | No authentication |
| SSL No Authentication | Neither the client nor the server authenticates itself to the other. No certificates are sent or exchanged. In this case, only SSL encryption / decryption is used. |
| SSL One Way | Only the directory server authenticates itself to the client. The directory server sends the client a certificate verifying that the server is authentic. |
| SSL Two Way | Both client and server authenticate themselves to each other. They do this by sending certificates to each other. |
| Simple | The client identifies itself to the server by means of a DN and a password which are sent in the clear over the network. The server verifies that the DN and password sent by the client matches the DN and password stored in the directory. |

   The bind mode is optional in subject specification. For the directive to be applicable, the bind mode specified on one node should match the bind mode specified on the node with which it is communicating.

**b.** Specify the entity or entities to whom you are granting access.

| Entity | Description |
|---|---|
| Everyone (*) | All who try to access the entry |
| A Specific Group | A previously defined group name |
| A Specific Entry | A previously defined directory entry |
| A Subtree | An entire subtree in the directory, which you select |
| When Session User's Distinguished Name (DN) Is Identified by Attribute | Anyone whose DN is an attribute in the entry. For example, you might want to grant read access to a group entry to members of the group |
| When Session User's Distinguished Name (DN) Matches the Accessed Entry | Anyone who has correctly logged in as the entry specified |

6. Select the Access Rights tab page.
   a. Determine what kinds of rights are granted: Browse, Add, Delete, or Unspecified. If an entry is unspecified, then access is determined at the next highest level in which access is specified.
   b. Click OK.

### Modifying Content Access Items of an Existing ACP by Using Oracle Directory Manager

If you configured Oracle Directory Manager to always display ACPs, as described in "Configuring the Display of ACPs in Oracle Directory Manager" on page 9-17, then you can navigate to and modify a content access item to an existing ACP as follows:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server* > Access Control Management. Select Access Control Management. All of the defined Access Control Policy Points (ACPs) appear in a list below Access Control Management in the navigator pane. They also appear in the right pane.
2. Under Access Control Management, select an ACP to display its information in the right pane, or double-click an ACP in the right pane to display the data in its own dialog box.

3. In the Content Access Items box, select the content access item you want to modify.
4. Just below the Content Access Item window, click Edit. The Content Access Items dialog box appears. Each tab page contains items you can modify.
5. Specify the items in the Entry Filter tab page (if applicable) as described in "Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-25.
6. Select the By Whom tab page and specify the items as described in the section "Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-25.
7. Select the Attribute tab page.
   a. From the right menu, select the attribute to which you want to grant or deny access.
   b. From the left menu, select the matching operation to be performed against the attribute. Choices are EQ (Equal (=)) and NEQ (Not Equal (!=)).
8. Select the Access Rights tab page and specify the items as described in the section "Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-25.
9. Click OK.

## Adding an ACP and Creating Access Items by Using Oracle Directory Manager

To create a new ACP:

1. In the navigator pane, expand, Oracle Internet Directory Servers > *directory server*. Select Access Control Management.
2. On the toolbar, click Create. A New Access Control Point dialog box appears.
3. In the Path to Entry field, enter the distinguished name (DN) of the entry that will be the ACP.

   **Note:** You can find the DN by looking in the navigator pane for the entry or by clicking Browse.

4. To define structural access items (entries), just below the Structural Access Items window, click Create. The Structural Access Item dialog box appears. Use

the tab pages in this dialog box as explained in the section "Modifying Structural Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-25.

5. To define content access items (attributes), just below the Content Access Items window, click Create. The Content Access Item dialog box appears. Use the tab pages in this dialog box as explained in the section "Modifying Content Access Items of an Existing ACP by Using Oracle Directory Manager" on page 9-28.
6. Click Ok to close this dialog box and return to the main Oracle Directory Manager dialog box.

## Example: Managing ACPs by Using Oracle Directory Manager

This example illustrates how to use Oracle Directory Manager to create a new ACP that has ACIs within it. Suppose you are an administrator in a large company, and you want to limit access to user passwords, so that everyone can compare a password, but only the owner of each password, that is, the user, can read the password or modify it.

In this example, we create a new ACP and populate it with four ACIs that set the following permissions:

- Limited access to a `userpassword` attribute by everyone
- Open access to the same `userpassword` attribute by the user himself
- Open access to all attributes except `userpassword` to everyone
- Open access to all attributes to everyone

### Create a New ACP

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server*, and select Access Control Management. A list of ACPs appears in the right pane.
2. Click Create at the bottom of the right pane. A New Access Control Point dialog box appears.
3. In the Path To Entry field, enter the DN where you want the ACP. The ACIs within the ACP will apply to all entries below and including that DN.

**Structural Access Items** To set the access rights for an entry:

1. Just below the Structural Access Items box, click Create. A Structural Access Items dialog box appears. It contains three tabs: Entry Filter, By Whom, and Access Rights.

   Because you want the ACIs to apply to all entries under the ACP, do not use the Entry Filter tab page.

2. Select the By Whom tab page to define the subject of the ACI. From the Bind Mode list, select the authentication mode appropriate to your environment. To create access rights for everyone, select Everyone. Click OK.

3. Select the Access Rights tab page. By default, all rights—browse, add, and delete—are granted.

   a. Change the access rights so that Everyone can browse all entries, but cannot add or delete them.

   b. Click Ok.

**Content Access Items** The four ACIs in this example use the same structural content item information. They differ only in the content access they allow. The rest of this section describes how to create the content access for the ACIs.

To define the content access items:

1. Below the Content Access Items box, click Create. The Content Access Items dialog box appears.

   Because you want this ACI to apply to all entries under the ACP, do not use the Entry Filter tab page.

2. Select the By Whom tab page, select Everyone, then click OK.

3. Select the Attribute tab page. This page has two fields. The first has two choices: EQ (equals) and NEQ (not equals). The second sets the attribute.

   Select EQ and select `userPassword`.

4. Select the Access Rights tab page. By default, all permissions are granted. Change the permissions so that read, search, write, and compare are denied.

5. Click Ok.

   You have completed one ACI.

### Create Another ACI

Create another ACI that allows a user to read, write, search, and compare his own password.

1. Under the Content Access Items box, click Create. The Content Access Items dialog box appears.
2. Select the By Whom tab page. Click When Session User's Distinguished Name (DN) Matches the Accessed Entry, then click OK.
3. Select the Attribute tab page. This tab page has two lists.The first has two choices: EQ (equals) and NEQ (not equals). The second sets the attribute.

   Select EQ and userPassword.
4. Select the Access Rights tab page.

   Grant access to read, search, write, and compare. Leave selfwrite unspecified.
5. Click Ok.

You have now created two ACPs. One denies Everyone read, search, write, and compare access to the `userPassword` attribute. The second allows the owner of the password to read, search, write, and compare that attribute.

### Create a Third ACI

The next ACI grants access to Everyone to read, search, and compare all attributes except `userPassword`. It denies write access.

1. Under the Content Access Items field, click Create to display the Content Access Items.
2. Select the By Whom tab page.

   Select Everyone, then click OK.
3. Select the Attribute tab page.

   Select NEQ and `userPassword`.

   This combination means that any attribute that is *not* equal to `userpassword` is the object of the permissions in this ACI.
4. Select the Access Rights tab page.

   Grant access to read, search, and compare. Deny write access. Leave selfwrite unspecified.
5. Click Ok to apply these permissions and close the dialog box.

### Create a Fourth ACI

The next ACI grants access to Self to read, browse, and write all attributes except `userpassword`. Including this ACI avoids any ambiguity about whether Self has the same access permissions as Everyone to attributes other than `userPassword`.

1. Under the Content Access Items field, click Create to display the Content Access Items dialog box.
2. Select the By Whom tab page.

   Click When Session User's Distinguished Name (DN) Matches the Accessed Entry. Click Ok.
3. Select the Attribute tab page.

   From the lists, select NEQ and `userPassword`. This combination means that any attribute that is *not* equal to `userPassword` is the object of the permissions in this ACI.
4. Press the Access Rights tab page.

   Grant access to read, search, and write. Leave Selfwrite unspecified.
5. Click Ok to apply these permissions and close the dialog box.

Consider other access restrictions you might want to implement. Your directory might contain many entries and attributes that should not be available to everyone.

## Granting Entry-Level Access by Using Oracle Directory Manager

To grant entry-level access by using Oracle Directory Manager:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory server* > Entry Management. You may either:
   - Select the entry to display its properties in the right pane
   - Use the search panel to find the entry, then double-click the entry to open the Entry dialog box.
2. Select the Local Access tab page, then create and edit local ACIs in the Structural Access Item and Content Access Item boxes.
3. Once you have made the changes, click Apply in the main Oracle Directory Manager window.

**Note:** You must click Apply to send the information you just entered to the directory server. If you do not click Apply, the information you just entered is simply held in the Oracle Directory Manager cache.

# Managing Access Control by Using Command Line Tools

As described in "Overview of Access Control Policy Administration" on page 9-2, directory access control policy information is represented as user modifiable operational attributes. Hence, you can manage directory access control by using the ldapmodify command to set and alter values of these attributes. Any tool, including ldapmodify and ldapmodifymt, can be used for this purpose.

In order to directly edit the ACI, you should understand the format and semantics of the directory representation of the ACI. This section contains the formal specification of the ACI format and a description of the semantic issues necessary to manage the ACI using command line tools.

**See Also:**

- "LDAP Data Interchange Format (LDIF) Syntax" on page A-2 for information about how to format input using **LDAP Data Interchange Format (LDIF)**, the required input format for line mode commands
- "ldapmodify Syntax" on page A-13 for information about how to run ldapmodify
- Appendix D, "Using Access Control Directive Format" for information about the format or syntax of ACI

## Examples: Managing Access Control

This section contains these topics:

- Example: Setting Up an Inheritable ACP by Using ldapmodify
- Example: Setting Up Entry-Level ACIs by Using ldapmodify
- Typical Access Control Policies

### Example: Setting Up an Inheritable ACP by Using ldapmodify

This example sets up subtree access permissions in an `orclACI` at the **Root DSE**. Because this example refers to the `orclACI` attribute, this access directive governs all the entries in the DIT.

Note the presence of the << EOF characters.

```
ldapmodify -v -h $1 -D "cn=Directory Manager, o=IMC, c=US" -w "controller" <<
EOF
dn:
changetype: modify
replace: orclaci
orclaci: access to entry
    by dn= "cn=directory manager, o=IMC, c=us" (browse, add, delete)
    by * (browse, noadd, nodelete)
orclaci: access to attr=(*)
    by dn= "cn=directory manager, o=IMC, c=us" (search, read, write, compare)
    by self (search, read, write, compare)
    by * (search, read, nowrite, nocompare)

EOF
```

### Example: Setting Up Entry-Level ACIs by Using ldapmodify

This example sets up entry-level access permissions in the `orclEntryLevelACI` attribute. Because this resides in the `orclEntryLevelACI` attribute, this ACL governs only the entry in which it resides. Note the presence of the << EOF characters.

```
ldapmodify -v -h $1 -D "cn=Directory Manager, o=IMC, c=US" -w "controller" <<
EOF
dn:
changetype: modify
replace: orclentrylevelaci
orclentrylevelaci: access to entry
    by dn= "cn=directory manager, o=IMC, c=us" (browse, add, delete)
    by * (browse, noadd, nodelete)
orclentrylevelaci: access to attr=(*)
    by dn= "cn=directory manager, o=IMC, c=us" (search, read, write, compare)
    by * (search, read, nowrite, nocompare)

EOF
```

---

**Note:** In this example, no DN value is specified. This means that this ACI pertains to the root DSE and its attributes only.

---

### Typical Access Control Policies

This section shows these advanced and typical examples of access control policies:

- Example: Using Wild Cards
- Example: Selecting Entries by DN
- Example: Using Attribute and Subject Selectors
- Example: Granting Read-Only Access
- Example: Granting Selfwrite Access to Group Entries

**Example: Using Wild Cards** This example shows the use of wild cards (*) in the object and subject specifiers. For all entries within the `acme.com` domain, it grants to everyone browse permission on all entries, as well as read and search permissions on all attributes.

`orclACI` attribute in the ACP at `dc=com`

```
access to entry by * (browse)
access to attr=(*) by * (search, read)
```

Note that, in order to allow reading the attributes, browse permissions must be granted on the entries in order for read permissions to be granted to the attributes of those entries.

**Example: Selecting Entries by DN** This example shows the use of a regular expression to select the entries by DN in two access directives. It grants to everyone under `dc=acme,dc=com` access to read address book attributes only. It extends only to everyone within `dc=us,dc=acme,dc=com` read access to all attributes.

`orclACI` attribute of `dc=acme, dc=com`:

```
access to entry by * (browse)
access to attr=(cn, telephone, email) by * (search, read)
```

`orclACI` attribute of `dc=us, dc=acme, dc=com`:

```
access to entry by * (browse)
access to attr=(*) by dn=".*,dc=us,dc=acme,dc=com" (search, read)
```

**Example: Using Attribute and Subject Selectors** This example shows the use of an attribute selector to grant access to a specific attribute, and various subject selectors. The example applies to entries in the `dc=us,dc=acme,dc=com` subtree. The policy enforced by this ACI can be described as follows:

- For all entries within the subtree, the administrator has add, delete, and browse permissions. Others within the `dc=us` subtree can browse, but those outside it have no access to the subtree.
- The salary attribute can be modified by one's manager and viewed by oneself. No one else has access to the salary attribute.
- The `userPassword` attribute can be viewed and modified by oneself and the administrator. Others can only compare this attribute.
- The `homePhone` attribute can be read and written by oneself and viewed by anyone else.
- For all other attributes, only the administrator can modify values. Everyone else can compare, search, read, but cannot update attribute values.

"orclACI" attribute of "dc=us, dc=acme, dc=com":

```
access to entry
by dn="cn=admin, dc=us,dc=acme,dc=com" (browse, add, delete)
by dn=".*, dc=us,dc=acme,dc=com" (browse)
by * (none)
```

```
access to attr=(salary)
by dnattr=(manager) (read, write)
by self (read)
by * (none)


access to attr=(userPassword)
by self (search, read, write)
by dn="cn=admin, dc=us,dc=acme,dc=com" (search, read, write)
by * (compare)


access to attr=(homePhone)
by self (search, read, write)
by * (read)


access to attr != (salary, userPassword, homePhone)
by dn="cn=admin, dc=us,dc=acme,dc=com" (compare, search, read, write)
by * (compare, search, read)
```

**Example: Granting Read-Only Access** This example gives to everyone under `dc=acme,dc=com` access only to read address book attributes. It also extends to everyone read access to all attributes within the `dc=us,dc=acme,dc=com` subtree only.

`orclACI` attribute of `dc=acme, dc=com`:

```
access to entry by * (browse)
access to attr=(cn, telephone, email) by * (search, read)
```

`orclACI` attribute of `dc=us, dc=acme, dc=com`:

```
access to entry by * (browse)
access to attr=(*) by dn=".*,dc=us,dc=acme,dc=com" (search, read)
```

**Example: Granting Selfwrite Access to Group Entries** This example allows people within the US domain to add or remove only their own name (DN) to or from the member attribute of a particular group entry, for example, a mailing list.

`orclEntryLevelACI` attribute of the group entry in question:

```
access to attr=(member)
by dn=".*, dc=us,dc=acme,dc=com" (selfwrite)
```

# 10

# Managing Directory Replication

Replication is the mechanism that maintains exact duplicates of specified naming contexts on multiple nodes.

> **Note:** For release 2.1.1, you can use Oracle Internet Directory replication only if you have installed **Advanced Symmetric Replication (ASR)**, which ships with all standalone purchases of Oracle Internet Directory and with Oracle8*i* Enterprise Edition. ASR is not included with Oracle8*i* Standard Edition.

This chapter contains these topics:



> **See Also:** "Distributed Directories: Replication" on page 2-26 for a conceptual discussion of replication

# Installing and Configuring Replication

This section describes how to install and initialize Oracle directory replication server software on a node.

Each node in a group of **DSA**s holds an updatable copy, also called an updatable replica, of the same set of **naming context**s. These naming contexts are synchronized with each other by replication processing. This group of nodes is called a **directory replication group (DRG)**.

---

**Note:** The instructions in this section apply to setting up replication in a group of empty nodes. For instructions on adding a node to an existing DRG, see "Adding a Replication Node" on page 10-19.

---

To install and configure a replication group, perform these general tasks:

- Task 1: Install Oracle Internet Directory on All Nodes in the DRG
- Task 2: Decide Which Node Will Serve as the ASR Master Definition Site (MDS)
- Task 3: At the MDS, Set Up ASR for a Directory Replication Group
- Task 4: Start Oracle Directory Server Instances on All the Nodes
- Task 5: Configure Replication
- Task 6: Start the Replication Servers on All the Nodes

---

**Note:** In Oracle Internet Directory release 2.1.1, procedures and tools are not available to create an environment (directory network) consisting of more than one DRG.

---

## Task 1: Install Oracle Internet Directory on All Nodes in the DRG

Note that the typical installation of the Oracle8*i* Enterprise Edition, which is required for the Oracle Internet Directory, includes Oracle **Advanced Symmetric Replication (ASR)**. By contrast, a typical installation of Oracle8*i* Standard Edition does not include ASR.

**See Also:** Installation documentation for Oracle Internet Directory

## Task 2: Decide Which Node Will Serve as the ASR Master Definition Site (MDS)

A **master definition site (MDS)** is any of the Oracle Internet Directory databases in which the administrator is going to run the configuration scripts. A remote master site is any site other than the Master Definition Site that participates in ASR replication.

You must be able to use **Net8** to connect to the MDS database and all other nodes that constitute the DRG.

## Task 3: At the MDS, Set Up ASR for a Directory Replication Group

The following sections lead you through installing and configuring ASR through Oracle Internet Directory installation scripts. More advanced ASR users may prefer to configure ASR through the Oracle8*i* Replication Manager Tool.

> **See Also:** Oracle8*i* Server replication documentation and the online help for Oracle8*i* Replication Manager Tool for information on configuring ASR with the Oracle8*i* Replication Manager Tool

Setting up the Oracle Advanced Symmetric Replication (ASR) environment to establish a Directory Replication Group (DRG) requires you to:

- Prepare the Net8 environment for replication
- Configure ASR for directory replication.

### Prepare the Net8 Environment for Replication

Follow these steps, described more fully below, on *all nodes* in the Directory Replication Group to prepare the Net8 environment:

1. Configure sqlnet.ora.
2. Configure tnsnames.ora.
3. Create rollback table space and rollback segments.
4. Modify the parameters in the initialization parameter file, init.ora.
5. Stop and restart the listener.
6. Stop and restart the Oracle Internet Directory database.

To prepare the Net8 environment for replication:

1. Configure `sqlnet.ora`.

   The `sqlnet.ora` file should contain the following parameters at minimum:

   ```
   names.directory_path = (TNSNAMES)
   names.default_domain = domain
   ```

On UNIX, this file is in $*ORACLE_HOME*/network/admin

On Windows NT, this file is in *ORACLE_HOME*\network\admin

2. Configure `tnsnames.ora`.

   The `tnsnames.ora` file must contain **connect descriptor** information in the following format for all Oracle Internet Directory databases:

   ```
   net_service_name =
       (DESCRIPTION =
         (ADDRESS =
            (PROTOCOL = TCP)
            (HOST = HOST_NAME_OR_IP_ADDRESS)
            (PORT = 1521))
       (CONNECT_DATA =
          (service_name = service_name))
   ```

On UNIX, this file is in $*ORACLE_HOME*/network/admin

On Windows NT, this file is in *ORACLE_HOME*\network\admin

---

**Note:** You may domain-qualify the net service name (for example, `sales.com`). Regardless of your choice, be sure that the domain component matches the one specified in the NAMES.DEFAULT_DOMAIN parameter in the `sqlnet.ora` file.

---

3. Create rollback table space and rollback segments.

   You may want to create multiple rollback segments. You can increase the size of the table spaces and segments to meet your system requirements.

   a. Create a tablespace for rollback segments.

      Execute SQL*Plus by typing the following command:

      ```
      sqlplus system/system_password@net_service_name
      ```

At the SQL*Plus prompt, type:

```
CREATE TABLESPACE table_space_name
datafile file_name_with_full_path SIZE 50M REUSE AUTOEXTEND ON NEXT
10M MAXSIZE max_bulk_update transaction_size ex:500M;
```

**b.** Create rollback segments.

At the SQL*Plus prompt, type the following lines for each rollback segment:

```
CREATE ROLLBACK SEGMENT rollback_segment_name
tablespace table_space_name storage (INITIAL 1M NEXT 1M OPTIMAL 2M
MAXEXTENTS UNLIMITED);
```

Repeat the `CREATE ROLLBACK SEGMENT` command for each rollback segment entered in the initialization parameter file.

**4.** Modify the parameters in the initialization parameter file, `init.ora`.

Type the following lines in the initialization parameter file:

```
rollback_segments = (rollback_segment_name_1, rollback_segment_name_2 ...)
JOB_QUEUE_PROCESSES = a_minimum_of_total_number_of_LDAP_nodes_minus_one
SHARED_POOL_SIZE = 20000000
OPEN_LINKS = a_minimum_of_total_number_of_LDAP_nodes_minus_one
```

---

**Note:** When setting the number of job queue processes, consider using a number high enough to accommodate any nodes you may want to add in the future.

---

Ensure that the total **System Global Area (SGA)** does not exceed 50% of your system's physical memory.

---

**Note:** Every time a database is started, a System Global Area (SGA) is allocated and Oracle background processes are started. The SGA is an area of memory used for database information shared by the database users. The combination of the background processes and memory buffers is called an Oracle instance.

---

**5.** Stop and restart the listener.

To stop the listener for the Oracle Internet Directory database, use the listener control utility (lsnrctl). Type the following command at the LSNRCTL command prompt:

```
SET PASSWORD password
STOP [listener_name]
```

`SET PASSWORD` is required only if the password is set in the `listener.ora` file. The password defaults to ORACLE. The default listener name is LISTENER.

To restart the listener for the Oracle Internet Directory database, type the following command at the LSNRCTL command prompt:

```
START [listener_name]
```

6. Stop and restart the Oracle Internet Directory database.

   To stop and restart the Oracle Internet Directory database, you can use SQL*Plus.

   **See Also:**

   - *Net8 Administrator's Guide*
   - *Oracle8i Administrator's Guide for* instructions on stopping and restarting the database

### Configure Oracle ASR For Directory Replication

To configure ASR for the replication group, complete the following steps *from the MDS*:

1. Log on as the Oracle Internet Directory software owner account from a UNIX prompt.
2. Change to the following directory:

- On UNIX: $*ORACLE_HOME*/ldap/bin
- On Windows NT: *ORACLE_HOME*\ldap\bin

**Note:** Before proceeding to the next step, connect as the system user on all nodes, including the MDS, from the MDS console. Ensure the following:

- The Oracle Internet Directory database is up and running
- The Oracle Internet Directory listener is up and running
- The connect descriptor is correct
- The system password is correct

3. Run the following script from the MDS:

```
ldaprepl.sh -asrsetup
```

This script executes a number of operations.

- It configures the MDS.
- It configures the remote master sites.
- It configures replication push jobs at all sites.
- It resumes replication at the MDS.
- It verifies that all steps have completed successfully.

As the script runs, it asks for the information in the following table, first for the MDS, then for the master sites.

| Information | Definition |
|---|---|
| Host name | Name of the computer |
| Global name | Net service name of the MDS database, as listed in the file `tnsnames.ora` |
| System password | system password |

After you have provided the necessary information for the first master site, the script asks if there is another master site.

4. Enter `Y` or `N`. If you enter `N`, to indicate that you have identified all sites, then it shows a table of the information you have provided, and asks for confirmation. If it is not correct, then press `N`. The script will start again at the beginning, asking about the MDS again.

After you have provided all the information, the script asks you to verify the correctness of the information. If the information is correct and you press Y, then the script begins configuring the sites.

This process may take a long time, depending on your system resources and the number of nodes in your DRG. The script keeps you informed of its progress.

---

**Note:** If you must interrupt the process before it is complete, then you must start at the beginning. Interrupting the process will not negatively affect your re-installation.

---

---

**Troubleshooting Tip:** If the process fails, then do the following:

1. Check the `$ORACLE_HOME/ldap/admin/logs/ldaprepl.log` file to see the status.
2. Go to the directory `$ORACLE_HOME/ldap/admin` and check the status of replication jobs by running the following command:

   ```
   sqlplus system/password@net_service_name @ldaplogq.sql
   ```

Run this command for each node in the DRG. Issuing this command should result in no rows being selected. If rows are selected containing the failed status and error messages, then this means that ASR set up failed. In this case, you may:

- Run the script from the beginning
- Consult the troubleshooting chapter in *Oracle8i Replication*
- Determine a solution from error message information by consulting an expert in Advanced Symmetric Replication (ASR)

---

---

**Note:** If you have large initial data requirements, then use the bulkload tool to load initial data on all the nodes in the DRG. You must stop the server before using bulkload, and bring it up again afterwards.

---

**See Also:**

- *Oracle8i Administrator's Guide* for instructions on ensuring that the database and listener are running
- *Net8 Administrator's Guide* for instructions on ensuring that the connect string is correct
- "bulkload Syntax" on page A-23 for bulkload syntax and usage notes

## Task 4: Start Oracle Directory Server Instances on All the Nodes

To start Oracle directory server instances on all nodes, run the following command:

```
oidctl connect=net_service_name server=oidldapd instance=instance_number_of_
ldap_server flags="-p port" start
```

**Note:** The `instance_number_of_ldap_server` need not be unique across the entire DRG. For example, you can have `instance=1` on both node A on node B.

**See Also:** Chapter 5, "Managing an Oracle Directory Server" for more information on starting an Oracle directory server **instance**

## Task 5: Configure Replication

You need to configure parameters for:

Oracle directory replication server

Oracle directory replication server configuration parameters are stored as special attributes in directory entries. You can configure replication parameters and replication agreements the same way you configure the Oracle Internet Directory. You can do either of the following:

- View and modify the agreements by using Oracle Directory Manager
- Alter the contents of the configuration entries and agreement entries through the command line tools, such as ldapadd and ldapmodify

This section explains both approaches.

| | |
|---|---|
| Replication agreements | Replication agreements are entries that list the member nodes within a replication group that share their changes. Replication agreements are referenced by Oracle directory replication server configuration parameters that load when the Oracle directory replication server runs. |

**Important:** When you install and configure replication for the first time, you must inform the Oracle directory replication server about the existence of the member nodes in the replication agreement. To do this, modify the `orclDirReplGroupDSAs` attribute in the replication agreement. See "Replication Agreement Parameters" on page 10-14 for more information.

### Location of Oracle Directory Replication Server Configuration Parameters

The Oracle directory replication server configuration parameters are stored in the replication server **configuration set entry**, which has the following DN:

```
cn=configset0,cn=osdrepld,cn=subconfigsubentry
```

This entry contains replication attributes that control replication processing. You can modify some of these attributes. Note that the `orclDirReplGroupAgreement` attribute contains a replication agreement identifier. In this release, only one replication agreement is possible.

### Oracle Directory Replication Server Parameters

The next table lists and describes the Oracle directory replication server configuration parameters.

| Parameter name | Description | Default Values | Modifiable? |
|---|---|---|---|
| modifyTimestamp | Time of entry creation or modification | | No |
| modifiersName | Name of person creating or modifying the entry | | No |
| orclChangeRetryCount | Single-valued attribute. The number of processing retry attempts for a change-entry before being dropped. The value for this parameter must be equal to or greater than 1 (one). | 10 | Yes |

| Parameter name | Description | Default Values | Modifiable? |
|---|---|---|---|
| orclPurgeSchedule | Single-valued attribute. Specifies purge (garbage collection) interval in minutes. Removes entries that are already applied or have been dropped as candidate changes. This thread is initiated periodically based on the frequency that you set. The value for this parameter must be equal to or greater than 1 (one). | 10 minutes | Yes |
| orclThreadsPerSupplier | Number of worker threads Oracle directory replication server provides for each supplier for change log processing. The value for this parameter must be equal to or greater than 1 (one). | 5 | Yes |
| orclDirReplGroupAgreement | Multi-valued attribute. Identifies the symmetrical replication agreements for which this server is responsible. | orclagreementid=000001, cn=orclreplagreements | No |
| orclChangeLogLife | Single-valued attribute. Specifies in hours the time for the life of entries in the change log store. 0 (zero) indicates that this is a change number-based purge.<br>**See Also:** "Change Log Purging" on page 2-30 | 0 | Yes |

### Viewing and Modifying Replication Configuration Parameters by Using Oracle Directory Manager

To view and modify replication configuration parameters:

1. In the navigator pane, expand Oracle Internet Directory > *directory_server_instance* > Server Management > Replication Server, then select the replication configuration set whose parameters you want to view or modify. The corresponding tab pages appear in the right pane.

   Configuration parameters appear in the General tab page. You can use this tab page to view replication configuration parameters, and modify many of them. The following table describes the fields in this tab page.

| Field | Description |
|---|---|
| Modify Timestamp | Time of entry creation or modification in **UTC (Coordinated Universal Time)**. You cannot modify this parameter. |
| Modifier's Name | Name of person creating or modifying the entry. You cannot modify this parameter. |
| Change Retry Count | Type the number of attempts that the conflict resolution process tries to apply each update before giving up and logging the incident. The default is 10. |
| Purge Schedule | Type the number of minutes in between garbage collections. The replication garbage collection thread removes entries that are already applied or have been dropped as candidate changes. The default is 10. |
| Number of Threads Per Supplier | Type the number of worker threads the Oracle directory replication server provides for each supplier for change log processing. The default is 5. |
| Set | Type the configuration identifier. |
| Change Log Life | Type the number of hours for the life of the change log objects.<br>**See Also:** "Change Log Purging" on page 2-30 |

## Modifying Replication Configuration Parameters by Using Command Line Tools

To modify replication configuration parameters by using command line tools, use the syntax documented in "ldapmodify Syntax" on page A-13.

**Modifying the Garbage Collection Interval by Using ldapmodify** This example uses an input file named `mod.ldif` to change the garbage collection interval from the default of 10 minutes to 30 minutes.

1. Edit `mod.ldif` as follows:

```
dn: cn=configset0,cn=osdrepld,cn=subconfigsubentry
changetype: modify
replace: orclPurgeSchedule
orclPurgeSchedule: 30
```

2. Use ldapmodify to update the replication server `configset0` parameter value as follows:

```
ldapmodify -h host -p port -f mod.ldif
```

3. Restart the Oracle directory replication server.

**Modifying the Change Log Life Parameter by Using ldapmodify** This example uses an input file named `mod.ldif` to change the change log life parameter to 10:

1. Edit `mod.ldif` as follows:

```
dn: cn=configset0,cn=oidrepld,cn=subconfigsubentry
changetype: modify
replace: orclChangeLogLife
orclChangeLogLife: 10 hours
```

2. Use ldapmodify to update the replication server `configset0` parameter value as follows:

```
ldapmodify -h host -p port -f mod.ldif
```

3. Restart the Oracle directory replication server.

**Modifying the Number of Retries Before a Change Is Moved into the Purge Queue by Using ldapmodify** This example uses an input file named `mod.ldif` to change the number of retry attempts from the default of ten times to five times. Specifically, after attempting to apply an update five times, the update is dropped and logged in the replication log.

1. Edit `mod.ldif` as follows:

```
dn: cn=configset0,cn=osdrepld,cn=subconfigsubentry
changetype: modify
replace: orclChangeRetryCount
orclChangeRetryCount: 5
```

2. Use ldapmodify to update the replication server `configset0` parameter value as follows:

```
ldapmodify -h host -p port -f mod.ldif
```

3. Restart the Oracle directory replication server.

**Modifying the Number of Worker Threads Used in Change Log Processing by Using ldapmodify** This example uses an input file named `mod.ldif` to change the number of worker threads used in change log processing:

1. Edit `mod.ldif` as follows:

```
dn: cn=configset0,cn=osdrepld,cn=subconfigsubentry
changetype: modify
replace: orclthreadspersupplier
orclthreadspersupplier: new_number_of_worker_threads
```

2. Use ldapmodify to update the replication server `configset0` parameter value as follows:

   ```
   ldapmodify -h host -p port -f mod.ldif
   ```

3. Restart the Oracle directory replication server.

   **See Also:** "Restarting Directory Server Instances" on page 3-7 for instructions on restarting the Oracle directory replication server

### Replication Agreement Parameters

In the parameter `DirectoryReplicationGroupDSAs`, type all of the host names of the DSAs in the DRG. Be sure that this information is identical on all the nodes.

**See Also:**

- "Viewing and Modifying Replication Agreement Parameters by Using Oracle Directory Manager" on page 10-15
- "Modifying Replication Agreement Parameters by Using ldapmodify" on page 10-16

### Location of Replication Agreement Parameters

Replication agreement parameters are stored in the replication agreement entries which have the following DN:

`orclAgreementID=`*`id number`*`,cn=orclreplagreements`

This entry contains attributes that pertain only to the nodes participating in this agreement. You can create multiple replication agreements to manage replication between reciprocating nodes, but you can reference only one of them in your start-server message by using Oracle Directory Manager. For Oracle Internet Directory release 2.1.1, only one replication agreement can be used.

The following table lists and describes the replication agreement parameters.

---

**Note:** Before you modify replication agreement parameters, be sure that you have started the Oracle Internet Directory on all nodes.

---

### Viewing and Modifying Replication Agreement Parameters by Using Oracle Directory Manager

To view and modify replication agreement parameters by using Oracle Directory Manager:

1. In the navigator pane, expand Oracle Internet Directory Servers > *directory_server_instance* > Server Management > Replication Server, and select Default Configuration Set.
2. In the right pane, select the Agreement tab to display the replication agreement.

   The fields in this tab page are described in the following table. You can view the parameters and modify some of them by double-clicking the attributes.

| Field in Oracle Directory Manager | Description | Default Values | Modifiable? |
|---|---|---|---|
| Agreements ID | Unique identifier for a replication agreement. | 000001 | No |
| Excluded Naming Contexts | Multi-valued attribute. Specifies naming contexts excluded from this replication agreement. Changes to entries in these naming contexts sent from other replicas are not applied on the local node. | None | Yes |
| Replication Group Nodes | Multi-valued attribute. Specifies nodes participating in symmetrical replication agreement. *Nodes that you specify here share updates with one another.* | | Yes |
| Update Schedule | Replication update interval for new changes and those being retried. The value is in minutes. | 1 | Yes |
| Orcl HIQSchedule | Replication update interval for the human intervention queue. The value is in minutes. The value is typically higher that orclUpdateSchedule. This gives administrators time to change the DIT structures when retrying an update fails to resolve a conflict. | 10 | Yes |
| Replication Protocol | Specifies the replication protocol used in this replication agreement. The supported protocol is ASR. | ODS_ASR_1.0 | No |

3. If you want to return to the values that appeared when you first opened this pane, then click Revert. If you are satisfied with your changes, then click Apply.

### Modifying Replication Agreement Parameters by Using ldapmodify

The following table lists and describes the replication agreement parameters.

| Parameter | Description | Default Values | Modifiable? |
|---|---|---|---|
| `orclAgreementID` | Unique identifier for a replication agreement. | 000001 | No |
| `orclExcludedNamingcontexts` | Multi-valued attribute. Specifies naming contexts excluded from this replication agreement. Changes to entries in these naming contexts sent from other replicas are not applied on the local node. | None | Yes |
| `orclDirReplGroupDSAs` | Multi-valued attribute. Specifies nodes participating in symmetrical replication agreement. *Nodes that you specify here share updates with one another.* | | Yes |
| `orclUpdateSchedule` | Replication update interval for new changes and those being retried. The value is in minutes. | 1 | Yes |
| `OrclHIQSchedule` | Replication update interval for the human intervention queue. The value is in minutes. The value is typically higher that orclUpdateSchedule. This gives administrators time to change the DIT structures when retrying an update fails to resolve a conflict. | 10 | Yes |

| Parameter | Description | Default Values | Modifiable? |
|---|---|---|---|
| orclReplicationProtocol | Specifies the replication protocol used in this replication agreement. The supported protocol is ASR. | ODS_ ASR_1.0 | No |

To add more nodes to the values in a replication agreement entry, run ldapmodify at the command line, referencing an LDIF-formatted file.

This example uses an input file named `mod.ldif` to add two nodes to a replication agreement:

1. Edit `mod.ldif` as follows:

```
dn: orclagreementid=000001,cn=orclreplagreements
changetype: modify
add: orcldirreplgroupdsas
orcldirreplgroupdsas: hollis
orcldirreplgroupdsas: eastsun-11
```

2. Use ldapmodify to update the replication server `configset0` parameter value as follows:

```
ldapmodify -h host -p port -f mod.ldif
```

3. Restart the Oracle directory replication server.

This procedure modifies the entry containing the replication agreement whose DN is `orclagreementid=000001,cn=orclreplagreements`. The input file adds the two nodes, hollis and eastsun-11, into the replication group governed by `oraclagreementid 000001`.

---

**Note:** You must include the new nodes—for example, hollis and eastsun-11 in the above sample LDIF file—in the `orclDirReplGroupDSAs` parameter on each node in the replicated environment before you start the replication process.

"Adding a Replication Node" on page 10-19 explains the process of adding a new node to a replication environment.

---

Because Oracle Internet Directory release 2.1.1 supports only one configuration set for Oracle directory replication server, you do not need to specify a configuration set.

## Task 6: Start the Replication Servers on All the Nodes

To start replication servers on all nodes, type the following command:

```
oidctl connect=db_connection_string server=oidrepld instance=1
   flags="-h host -p port" start
```

Note that the instance number does not need to be unique across the entire DRG.

> **See Also:** Chapter 5, "Managing an Oracle Directory Server" for information on starting the replication servers

### Using the Change Log Flag

You can turn off change logging, which occurs in the Oracle directory server, by using the default value of the `-l` flag in the OID Control Utility command for Oracle directory server from *true* to *false*. This is useful if you suspect that the change log file might not be emptying. However, turning change logging off on a given node means that updates on that node cannot be replicated to other nodes in the DRG.

### Using the Multimaster Flag

You can turn off the multimaster flag, which occurs in the Oracle directory replication server, by using the default value of the `-m` flag in the OID Control Utility command for Oracle directory server from *true* to *false*. This is useful for reducing performance overhead if you are deploying a single master with read-only replica consumers. The multimaster option controls conflict resolution, which serves no purpose if you are deploying a single master.

> **See Also:** "Conflict Resolution in Replication" on page 2-31

# Adding a Replication Node

There are two ways to add a new node to a live replication group.

- Using ldifwrite

  This method, described in this section, is the easier of the two. The process can be fully automated, and the generated file can be used for partial replication. Use this procedure unless your directory is very large. Backup using this method can take up to seven hours for a directory with one million entries.

- Using cold backup

  This method, described in Appendix B, "Adding a DSA Using the Database Copy Procedure", cannot be fully automated and cannot be reused for partial replication. However, cold backup takes much less time for a large directory server. If your directory has, say, more than a million entries, then use this method.

---

**Note:** Before you add a replication node, prepare the Net8 environment. For instructions, see "Prepare the Net8 Environment for Replication" on page 10-3.

---

To add a replication node to a functioning DRG of any significant size, follow these steps, each of which is more fully described later in this chapter.

Task 1: Stop the Oracle Directory Replication Server on All Nodes

Task 2: Configure the New Node into the LDAP Replication Group on All the Existing Nodes

Task 3: Identify a Sponsor Node and Switch the Sponsor Node to Read-Only Mode

Task 4: Backup the Sponsor Node by Using ldifwrite

Task 5: Perform ASR Add Node Setup

Task 6: Switch the Sponsor Node to Updatable Mode

Task 7: Start the Oracle Directory Replication Server on All Nodes Except the New Node

Task 8: Load Data into the New Node by Using bulkload

Task 9: Start LDAP Server on the New Node

Task 10: Configure the LDAP Replication Agreement on the New Node

Task 11: Start the Oracle Directory Replication Server on the New Node

**Note:** Commands shown in the following steps require that the following variables be stored in the corresponding directories:

- Binaries: $*ORACLE_HOME*/bin
- SQL scripts: $*ORACLE_HOME*/ldap/admin
- UNIX scripts: $*ORACLE_HOME*/ldap/bin

Before beginning Task 1, be sure that all three of these variables are in the path.

## Task 1: Stop the Oracle Directory Replication Server on All Nodes

To stop the Oracle directory replication server, run the following command on each node in the LDAP replication group:

```
oidctl connect=db_connect_string server=oidrepld instance=1 stop
```

**Note:** The instance number may not be 1. Check the running process to discover the instance number in use here.

## Task 2: Configure the New Node into the LDAP Replication Group on All the Existing Nodes

The following example creates an LDIF file, add_node.ldif, and configures it into the replication group on all the existing nodes.

```
dn: orclagreementid=000001,cn=orclreplagreements
changetype: modify
replace: orcldirreplgroupdsas
orcldirreplgroupdsas: host_name_of_the_new_node
orcldirreplgroupdsas: host_name_of_existing_node_1
orcldirreplgroupdsas: host_name_of_existing_node_2
.
.
.
orcldirreplgroupdsas: host_name_of_existing_node_n
```

Run the following command against each node in the LDAP replication group:

```
ldapmodify -h host_name_of_the_node -p port -f add_node.ldif
```

**Note:** This command can be run from one work station for all nodes.

## Task 3: Identify a Sponsor Node and Switch the Sponsor Node to Read-Only Mode

A sponsor node is one that will supply the data to the new node. To identify a sponsor node and switch it to read-only mode:

1. Create a new file, `change_mode.ldif`, containing the following:

   ```
   dn:
   changetype: modify
   replace: orclservermode
   orclservermode: r
   ```

2. Run the following commands against the identified sponsor node:

   ```
   ldapmodify -D "cn=orcladmin" -w welcome -h host_name_of_sponsor_node
   -p port -f change_mode.ldif

   oidctl connect=net_service_name server=oidldapd restart
   ```

This restarts all running Oracle directory servers on the sponsor node in Read-Only mode. It takes approximately fifteen seconds for a directory server to restart.

**Note:** While the sponsor node is in read-only mode, you may not make any updates to it. You may, however, update any of the other nodes, but those updates are not replicated immediately.

Also, the sponsor node and the **MDS** may be the same node.

## Task 4: Backup the Sponsor Node by Using ldifwrite

Because this may take a long time, you may start "Task 5: Perform ASR Add Node Setup" while backup is in process.

Enter the following command:

```
ldifwrite -c db_connect_string -b "" -f output_ldif_file
```

## Task 5: Perform ASR Add Node Setup

You can perform this task at the same time as you are performing "Task 4: Backup the Sponsor Node by Using ldifwrite".

From the sponsor node, run the following script:

```
ldaprepl.sh -addnode
```

This script executes a number of operations.

- It quiesces ASR at the sponsor node and other existing **master sites**.
- It configures the master sites and the new node. A master site is any site other than the sponsor node that participates in LDAP replication.
- It configures replication push jobs at all sites including the new node.
- It checks that all steps have completed successfully. (This may take a long time.)
- It performs post-add-node operation.

As the script runs, it asks for the information in Table 10–1, first for the sponsor node then for the existing master sites.

***Table 10–1 ASR Setup Information***

| Information | Description |
|---|---|
| Host Name of sponsor node | Name of the computer |
| Global name | Net service name of the MDS or master site database, as listed in `tnsnames.ora` |
| system password | system password |

When you have identified all the existing master sites, enter `N`. The script then asks for information regarding the new node. Once you have provided that information, the script shows you a table of the information you have provided, and asks for confirmation.

If the information is not correct, then press `N`. The script then starts again at the beginning, asking the same information. If the information is correct and you enter `Y`, then the script begins configuring the sites.

This process can take a long time, depending on your system resources and the size of your DRG. The script keeps you informed of its progress.

**Note:** If for any reason you must interrupt the process before it is complete, then you must start from the beginning.

**Troubleshooting Tip:** If the process fails, then do the following:

1. Check the $*ORACLE_HOME*/ldap/admin/logs/ldaprepl.log file to see the status.
2. Go to the directory $*ORACLE_HOME*/ldap/admin and check the status of replication jobs by running the following command:

   ```
   sqlplus system/password@net_service_name @ldaplogq.sql
   ```

Run this command for each node in the DRG. Issuing this command should result in no rows being selected. If rows are selected containing the status [failed] and error messages, then this means that ASR set up failed. In this case, you may:

- Run the script from the beginning
- Consult the troubleshooting chapter in *Oracle8i Replication*
- Determine a solution from error message information by consulting an expert in Advanced Symmetric Replication (ASR)

## Task 6: Switch the Sponsor Node to Updatable Mode

To switch the sponsor node to updatable mode:

**1.** Edit change_mode.ldif to the following:

```
dn:
changetype: modify
replace: orclservermode
orclservermode: rw
```

**2.** Run the following commands on the sponsor node:

```
ldapmodify -D "cn=orcladmin" -w welcome -h host_name_of_sponsor_node
-p  port  -f change_mode.ldif

oidctl connect=net_service_name server=oidldapd restart
```

**Note:** Task 6 is very similar to Task 3. The only difference is that the `orclservermode` parameter in `change_mode.ldif` is being set back to `rw`, that is, Read-Write, in this step.

## Task 7: Start the Oracle Directory Replication Server on All Nodes Except the New Node

To start the Oracle directory replication server, type the following command:

```
oidctl connect=db_connection_string server=oidrepld instance=1
flags="-h host -p port" start
```

Verify that no directory or replication processes are running on the new node.

## Task 8: Load Data into the New Node by Using bulkload

To load data, type the following command:

```
bulkload.sh -connect db_connect_string_of_new_node -generate -load
-restore absolute_path_to_the_ldif_file_generated_by_ldifwrite
```

## Task 9: Start LDAP Server on the New Node

To start the LDAP server, type the following command:

```
oidctl connect=db_connect_string_of_new_node server=oidldapd
instance=1 flags="-p port" start
```

## Task 10: Configure the LDAP Replication Agreement on the New Node

Run the following command against the new node:

```
ldapmodify -h host_name_of_the_new_node -p port -f add_node.ldif
```

## Task 11: Start the Oracle Directory Replication Server on the New Node

To start the Oracle directory replication server, type the following command:

```
oidctl connect=db_connect_string_of_new_node server=oidrepld instance=1
flags="-h host_name_of_new_node -p port" start
```

# Deleting a Replication Node

You can delete a replication node from a **DRG** only if there are more than two nodes in the DRG.

To delete a replication node from a directory with fewer than a million entries, follow these steps, each of which is more fully described later.

Task 1: Stop the Oracle Directory Replication Server on All Nodes

Task 2: Stop All Processes in the Node to be Deleted

Task 3: Delete the Node from the Master Definition Site

Task 4: Start the Oracle Directory Replication Server on All Nodes

Task 5: Delete the Node from the Replication Group

Task 6: Restart the Oracle Directory Replication Server on the Remaining Nodes

---

**Note:** Commands shown in the following steps require that the following variables be stored in the corresponding directories:

- Binaries: $*ORACLE_HOME*/bin
- SQL scripts: $*ORACLE_HOME*/ldap/admin
- UNIX scripts: $*ORACLE_HOME*/ldap/bin

Before beginning Task 1, be sure that all three variables are in the path.

---

## Task 1: Stop the Oracle Directory Replication Server on All Nodes

To stop the Oracle directory replication server, run the following command on each node in the DRG:

```
oidctl connect=net_service_name server=oidrepld instance=1 stop
```

---

**Note:** The instance number may vary.

---

## Task 2: Stop All Processes in the Node to be Deleted

Stop the **OID Control Utility** and the **OID Monitor**.

**See Also:**

- "Stopping an Oracle Directory Server Instance" on page 3-5 for instructions about stopping the OID Control Utility
- "Stopping the OID Monitor" on page 3-3 for instructions about stopping the OID Monitor

## Task 3: Delete the Node from the Master Definition Site

From the **MDS**, run the following script:

```
ldaprepl.sh -delnode
```

This script executes these operations:

- It quiesces **ASR** at the MDS and other existing **master sites**.
- It deletes the node from the `orclDirReplGroupDSAs` parameter.
- It verifies that all steps have completed successfully.

As the script runs, it asks for the information in Table 10–2, first for the Master Definition Site then for the node to be deleted.

***Table 10–2 ASR Setup Information***

| Information | Description |
|---|---|
| Host Name of MDS or master site | Name of the computer |
| Global name | Net service name of the MDS or master site database, as listed in `tnsnames.ora` |

Once you have provided that information, the script shows you a table of the information you have provided, and asks for confirmation. If the information is not correct, then press `N`. The script then starts again at the beginning, asking the same information. If the information is correct and you enter `Y`, then the script begins configuring the sites.

This process can take a long time, depending on your system resources and the size of your DRG. The script keeps you informed of its progress.

**Note:** If, for any reason, you must interrupt the process before it is complete, then you must start from the beginning.

**Troubleshooting Tip:** If the process fails, then do the following:

1. Check the $*ORACLE_HOME*/ldap/admin/logs/ldaprepl.log file to see the status.
2. Go to the directory $*ORACLE_HOME*/ldap/admin and check the status of replication jobs by running the following command:

   ```
   sqlplus system/password@net_service_name @ldaplogq.sql
   ```

Run this command for each node in the DRG. Issuing this command should result in no rows being selected. If rows are selected containing the status [failed] and error messages, then this means that ASR set up failed. In this case, you may:

- Run the script from the beginning
- Consult the troubleshooting chapter in *Oracle8i Replication,*
- Determine a solution from error message information by consulting an expert in Advanced Symmetric Replication (ASR)

## Task 4: Start the Oracle Directory Replication Server on All Nodes

To start the Oracle directory replication server, type the following command:

```
oidctl connect=net_service_name server=oidrepld instance=1
flags="-h host -p port" start
```

## Task 5: Delete the Node from the Replication Group

Before deleting the node from the replication group, be sure that all of its changes have been applied to the other nodes.

The following example creates an LDIF file, `delete_node.ldif`, and configures it into the replication group on all the existing nodes. Notice that this LDIF file does not include the host name of the node to be deleted.

```
dn: orclagreementid=000001,cn=orclreplagreements
changetype: modify
replace: orcldirreplgroupdsas
orcldirreplgroupdsas: host_name_of_existing_node1
orcldirreplgroupdsas: host_name_of_existing_node2
.
.
.
orcldirreplgroupdsas: host_name_of_existing_node_n
```

Run the following command against each node in the LDAP replication group:

```
ldapmodify -h host_name_of_the_node -p port -f delete_node.ldif
```

## Task 6: Restart the Oracle Directory Replication Server on the Remaining Nodes

After deleting the node, restart the Oracle directory replication server on the remaining nodes for greater efficiency. To do this, type the following command:

```
oidctl connect=db_connection_string server=oidrepld instance=1
flags="-h host -p port" restart
```

# Resolving Conflicts Manually

This section contains these topics:



## Monitoring Replication Change Conflicts

If a conflict has been written into the log, then it means that the system is not able to resolve it by following its resolution procedure. To avoid further replication change conflicts arising from earlier unapplied changes, it is important to monitor the logs regularly.

To monitor replication change conflicts, examine the contents of the replication log. You can distinguish between messages by their respective timestamps.

## Examples of Conflict Resolution Messages

Conflict resolution messages, examples of which are shown below, are logged in the file `oidrepld00.log`. The path for this file is *`ORACLE_HOME`*`/ldap/log`. The result of each attempt to resolve the replication conflict is displayed at the end of each conflict resolution message.

### Example 1: An Attempt to Modify a Non-Existent Entry

```
2000/08/03::10:59:05:  ************ Conflict Resolution Message ************
2000/08/03::10:59:05:  Conflict reason: Attempted to modify a non-existent
entry.
2000/08/03::10:59:05:  Change number:1306.
2000/08/03::10:59:05:  Supplier:eastlab-sun.
2000/08/03::10:59:05:  Change type:Modify.
2000/08/03::10:59:05:  Target
DN:cn=ccc,ou=Recruiting,ou=HR,ou=Americas,o=IMC,c=US.
2000/08/03::10:59:05:  Result: Change moved to low priority queue after failing
on 10th retry.
```

### Example 2: An Attempt to Add an Existing Entry

```
2000/08/03::10:59:05:  ************ Conflict Resolution Message ************
2000/08/03::10:59:05:  Conflict reason: Attempted to add an existing entry.
2000/08/03::10:59:05:  Change number:1209.
2000/08/03::10:59:05:  Supplier:eastlab-sun.
2000/08/03::10:59:05:  Change type:Add.
2000/08/03::10:59:05:  Target DN:cn=Lou Smith, ou=Recruiting, ou=HR,
ou=Americas, o=IMC, c=US.
2000/08/03::10:59:05:  Result: Deleted duplicated target entry which was created
later than the change entry. Apply the change entry again.
```

### Example 3: An Attempt to Delete a Non-Existent Entry

```
2000/08/03::10:59:06:  ************ Conflict Resolution Message ************
2000/08/03::10:59:06:  Conflict reason: Attempted to delete a non-existent
entry.
2000/08/03::10:59:06:  Change number:1365.
2000/08/03::10:59:06:  Supplier:eastlab-sun.
2000/08/03::10:59:06:  Change type:Delete.
2000/08/03::10:59:06:  Target DN:cn=Lou
Smith,ou=recruiting,ou=hr,ou=americas,o=imc,c=us.
2000/08/03::10:59:06:  Result: Change moved to low priority queue after failing
on 10th retry.
```

## Using the Human Intervention Queue Manipulation Tool

The human intervention queue manipulation tool enables you to move the changes from the human intervention queue to either the retry queue or the purge queue. Moving the change to the purge queue means that there are no further attempts to re-apply the changelog entry. Perform the following general steps to address changes in the human intervention queue:

1. Shutdown the Oracle directory replication server.
2. Analyze the replication log.
3. Use the human intervention queue manipulation tool to move the changes to either the retry queue or the purge queue as described in the following sections.

### Moving a Change from the Human Intervention Queue into the Retry Queue

To place a change back into the retry queue, use this syntax:

```
hiqretry.sh –connect net_service_name [-start change_number]
[-end change_number] [-equal change_number] -supplier supplier_node
```

The arguments are:

| Argument | Description |
|---|---|
| -connect *net_service_name* | Connects to the database using the net service name defined in the `tnsnames.ora` file |
| -start *change_number* | Specifies the start change number for the retry operation. If you skip this option, then the command moves all the changes with change numbers less than or equal to the specified end change number back to the retry queue. |
| -end *change_number* | Specifies the end change number for the retry operation. If you skip this option, then the command moves all the changes with change numbers greater than or equal to the specified start change number back to the retry queue. |
| -equal *change_number* | Specifies the change number. The command moves the exact change conflict back to the retry queue. This option should not be present when `-start` or `-end` is used. |
| -supplier *supplier_node* | Specifies the supplier node where the changes originate |

### Moving a Change from the Human Intervention Queue into the Purge Queue

To place a change into the purge queue, use this syntax:

```
hiqpurge.sh –connect net_service_name [-start change_number] [-end change_
number] [-equal change_number] -supplier supplier_node
```

Arguments are:

| Argument | Description |
|---|---|
| -connect *net_service_name* | Connects to the database using the net service name defined in the tnsnames.ora file |
| -start *change_number* | Specifies the start change number for the purge operation. If you skip this option, then the command moves all the changes with change numbers less or equal to the specified end change number back to the purge queue. |

| Argument | Description |
|---|---|
| -end *change_number* | Specifies the end change number for the purge operation. If you skip this option, then the command moves all the changes with change numbers greater or equal to the specified start change number back to the purge queue. |
| -equal *change_number* | Specifies the change number of the change. The command moves the exact change conflict back to the purge queue. This option should not be present when `-start` or `-end` is used. |
| -supplier *supplier_node* | Specifies the supplier node where the changes originate |

**Note:** When using hiqretry.sh or hiqpurge.sh, if you do not want all changes to be moved, then you must supply either the `-equal` flag, or a combination of the `-start` and `-end` flags.

## Examples: Using the Human Intervention Queue Manipulation Tool

The following examples illustrate how to use the human intervention queue manipulation tool.

**Example: Retrying and Discarding Changes** Suppose that, after analyzing the replication log, you decide to do the following:

- Retry changes coming from the supplier node, ldap_rep1, with change numbers between 10324 to 10579
- Discard changes with change numbers between 10581 to 10623.

To do this, you issue these two commands:

```
hiqretry.sh –connect oiddb1 -start 10324 -end 10579 –supplier ldap_rep1
hiqpurge.sh –connect oiddb1 -start 10581 -end 10623 –supplier ldap_repl
```

The first command moves changes originating in ldap_rep1 with change numbers from 10324 to 10579 back to the retry queue. The second command deletes changes that originate in the supplier ldap_repl and that have change numbers from 10581 to 10623.

**Example: Moving a Single Change from the Human Intervention Queue to the Retry Queue** The following command moves the change with change number equal to 10519 back to the retry queue.

```
hiqretry.sh –connect oiddb1 –equal 10519 –supplier ldap_repl
```

**Example: Moving a Group of Changes from the Human Intervention Queue to the Retry Queue** The following command moves all the changes with change number greater or equal to 10324 back to the retry queue.

```
hiqretry.sh –connect oiddb1 –start 10324 –supplier ldap_repl
```

The following command moves all the changes with change numbers less than or equal to 10579 back to the retry queue.

```
hiqretry.sh –connect oiddb1 –end 10579 –supplier ldap_repl
```

**Example: Moving All Changes from the Human Intervention Queue to the Retry Queue** The following command includes no options. It moves all changes that originate in the supplier ldap_repl from the Human Intervention Queue to the retry queue.

```
hiqretry.sh –connect oiddb1 –supplier ldap_repl
```

## Using the OID Reconciliation Tool

When the Oracle directory replication server encounters inconsistent data, you can use the OID reconciliation tool to synchronize the entries on the consumer with those on the supplier. When you do this, perform the following general steps:

1. Set the supplier and the consumer to read-only mode.
2. Ensure that the supplier and the consumer are in tranquil state. If they are not in a tranquil state, then wait until they have finished updating.
3. Identify the inconsistent entries or subtree on the consumer.
4. Use the OID reconciliation tool to fix the inconsistent entries or subtree on the consumer.
5. Set the participating supplier and consumer back to read-write mode.

### Reconciling Inconsistent Data by Using the OID Reconciliation Tool

The OID reconciliation tool uses this syntax:

```
oidreconcile -h supplier_host -c consumer_host [-P supplier_port] [-p consumer_
port] [-s scope] -b basedn -W supplier_password -w consumer_password [-T thread]
```

| Argument | Description |
|---|---|
| -h *supplier_host* | Supplier host. This can be a computer name or IP address. |
| -c *consumer_host* | Consumer host. This can be a computer name or IP address. |
| -P *supplier_port* | Supplier TCP port. If you do not specify this option, then the tool connects to the default port (389). |
| -p *consumer_port* | Consumer TCP port. If you do not specify this option, then the tool connects to the default port (389). |
| -s *scope* | Reconcile scope: subtree |
| -b *basedn* | Specifies the distinguished name of the entry on which to perform reconciliation. |
| -W *supplier_password* | The password of `cn=orcladmin` of the supplier node |
| -w *consumer_password* | The password of `cn=orcladmin` of the consumer node |
| -T *thread* | Worker thread |

### How the OID Reconciliation Tool Works

When the OID reconciliation tool receives the specified DN, it compares the `orclGuid` of the parent DN on both the supplier and the consumer.

If the global identification (`orclGuid`) of both parents match, and the option `-s subtree` is set, then the OID reconciliation tool does the following:

1. Deletes all the entries in the subtree on the consumer node
2. Replaces them with entries from the supplier node

For example, the following command replaces the whole subtree starting from `"ou=hr,o=acme,c=us"` on the consumer with the equivalent subtree on the supplier:

```
oidreconcile -h supplier_host -P 389 -c consumer_host -p 389 -b
"ou=hr,o=acme,c=us" -s subtree -W supplier_password -w consumer_password
```

If the global identification (`orclGuid`) of both parents (`"o=acme,c=us"`) match, and `-s subtree` is not set, then the OID reconciliation tool replaces only the entry itself on the consumer node with the specified entry from the supplier node.

For example, the following command, in which the option `"-s subtree"` is not set, replaces only the specified entry, `"ou=hr,o=acme,c=us"`.

```
oidreconcile -h supplier -P 389 -c consumer -p 389 -b "ou=hr, o=acme, c=us"
-W supplier_password -w consumer_password
```

Figure 10–1 helps to explain how this process works.

***Figure 10–1 Example: OID Reconciliation Tool Process***

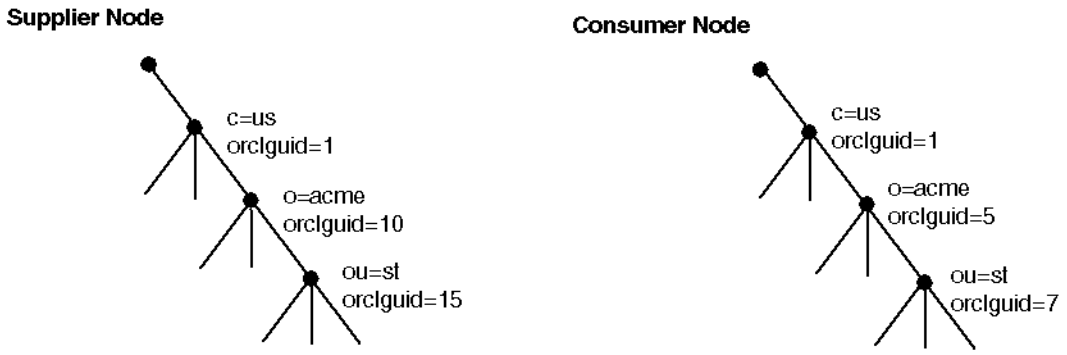


This figure shows two **DITs**, one on a supplier node and one on a consumer node. In the DIT on the supplier node, the `orclGuid` for c=us is 1 (one), the `orclGuid` for o=acme is 10, and the `orclGuid` for ou=st is 15. On the consumer node, the `orclGuid` for o=acme is 5, and the `orclGuid` for ou=st is 7.

The `orclGuid`s for the parent of `o=acme,c=us`—namely, `c=us`—on both the supplier and the consumer match. Therefore, the following command replaces all entries under `o=acme,c=us` on the consumer with the corresponding ones on supplier:

```
oidreconcile -h supplier -c consumer -b "o=acme, c=us" -s subtree -W supplier_
password -w consumer_password
```

If the `orclGuid` of both parents does not match, then the OID reconciliation tool does not perform the reconciliation. Instead, it tells the user the first ancestor on the consumer in which the `orclGuid` matches that of the same ancestor on the supplier.

For example, in Figure 10–1, suppose you were to run the following command:

```
oidreconcile -h supplier -c consumer -b "ou=st, o=acme, c=us" -s subtree
-W supplier_password -w consumer_password
```

As a result of this command, you would receive a message that the first ancestor of `ou=st` in which the `orclGuid`s `match` is `o=acme,c=us`. This message means that you should use `o=acme,c=us` as `basedn` argument to oidreconcile.

# 11

# Synchronizing with Multiple Directories

Oracle Internet Directory release 2.1.1 enables synchronization with supported third party metadirectory solutions. Synchronization with these metadirectory solutions occurs through the use of change logs. This chapter describes how that change log information is generated and used by supporting solutions. It also provides instructions for enabling other directories to synchronize with Oracle Internet Directory.

This chapter contains these topics:

# The Synchronization Process

Changes in an Oracle Internet Directory are recorded as entries in the change log object store. Other directories must have access to that store if they are to synchronize with Oracle Internet Directory. You grant them this access by registering them with Oracle Internet Directory.

Each entry in the change log store has a change number. Another directory retrieves from Oracle Internet Directory only those entries with change numbers equal to or greater than the last change it retrieved. For example, suppose that the entry that a directory last retrieved had a change number of 250. Entries that this directory subsequently retrieves must have change numbers of 250 or greater.

---

**Note:** If the entry with the change number matching the last change retrieved is not returned in the search results, then it means that some of the entries in the Oracle Internet Directory change log have been purged. The directory must then read the entire Oracle Internet Directory change log to synchronize its copy with that of Oracle Internet Directory.

---

Once you have registered another directory with Oracle Internet Directory, that directory can authenticate to Oracle Internet Directory and retrieve updates from it. It does this by following the processes described in this section.

> **See Also:** Enabling Other Directories to Synchronize with Oracle Internet Directory on page 11-4 for instructions on registering directories with Oracle Internet Directory

This section contains these topics:

- How a Directory Retrieves Changes the First Time from Oracle Internet Directory
- How a Connected Directory Updates the orclLastAppliedChangeNumber Attribute in Oracle Internet Directory
- How a Directory Retrieves Changes After the First Time from Oracle Internet Directory

## How a Directory Retrieves Changes the First Time from Oracle Internet Directory

In this example, `my_other_directory` acquires changes from Oracle Internet Directory by issuing the following command through ldapsearch:

```
ldapsearch -h host -p port -b "cn=changeLog" -s one
(&(objectclass=changeLogEntry)
(changeNumber >= orclLastAppliedChangeNumber )
( ! (modifiersname =cn=my_other_directory,cn=Subscriber Profile,
     cn=ChangeLog Subscriber,cn=Oracle Internet Directory ) ) )
```

When the directory is retrieving changes for the first time, the value for `orclLastAppliedChangeNumber` is the number you set in "Task 2: Register a Directory as a Change Subscription Object in Oracle Internet Directory" on page 11-5.

The argument `(!(modifiersname=client_bind_dn))` in the filter ensures that Oracle Internet Directory does not return changes made by the other directory itself.

## How a Connected Directory Updates the orclLastAppliedChangeNumber Attribute in Oracle Internet Directory

After retrieving changes from Oracle Internet Directory, the connected directory updates the `orclLastAppliedChangeNumber` attribute in its change subscription object. This allows Oracle Internet Directory to purge changes that connected directories have already applied. It also enables the connected directory to retrieve only the most recent changes, ignoring those it has already applied.

This example uses an input file named `mod.ldif` in which the last applied change number is 121. The connected directory updates `orclLastAppliedChangeNumber` in its change subscription object as follows:

1. Edit `mod.ldif`:

   ```
   dn: cn=my_other_directory,cn=Subscriber Profile,
       cn=ChangeLog Subscriber,cn=Oracle Internet Directory
   changetype:modify
   replace: orclLastAppliedChangeNumber
   orclLastAppliedChangeNumber: 121
   ```

2. Use ldapmodify to load the edited `mod.ldif` file:

   ```
   ldapmodify -h host -p port -f mod.ldif
   ```

**See Also:** "Change Log Purging" on page 2-30 for information about purging changes according to change numbers.

### How a Directory Retrieves Changes After the First Time from Oracle Internet Directory

To retrieve changes after the first time, the other directory issues a command by using ldapsearch. The following example returns all the changes with `changeNumber` equal to or greater than 121, except those related to operations performed by the other directory itself.

```
ldapsearch -h my_host> -p my_port_number -b "cn=changeLog" -s one"
(&(objectclass=changeLogEntry)  (changeNumber >=  122 )
( ! (modifiersname = cn=my_other_directory,cn=Subscriber Profile,
     cn=ChangeLog Subscriber,cn=Oracle Internet Directory ) ) )
```

## Enabling Other Directories to Synchronize with Oracle Internet Directory

To enable other directories to retrieve the changes stored in Oracle Internet Directory, you perform the tasks described in this section. This section contains these topics:

- Task 1: Perform Initial Bootstrapping
- Task 2: Register a Directory as a Change Subscription Object in Oracle Internet Directory
- Task 3: Grant Directories Access to the Oracle Internet Directory Change Log Object Store

### Task 1: Perform Initial Bootstrapping

To bootstrap a directory to synchronize data between a local directory and Oracle Internet Directory, execute these steps:

1. Retrieve the current change number from Oracle Internet Directory by executing the following command:

   `oidcurrentchange.sh -connect` *`net_service_name`*

   This displays the current change number. Later, you will use this number to fill the `orclLastAppliedChangeNumber` field when you register the directory.

2. Use ldifwrite to export data from Oracle Internet Directory into an LDIF file.
3. Convert the LDIF file to a format suitable to the client directory, then load it into the client directory.

---

**Note:** Initial bootstrapping is not required with a newly installed Oracle Internet Directory. In this case, the current change number of the newly installed Oracle Internet Directory is 0 (zero).

---

**See Also:** "ldifwrite Syntax" on page A-27 for instructions on using ldifwrite

## Task 2: Register a Directory as a Change Subscription Object in Oracle Internet Directory

To enable other directories to synchronize with an Oracle Internet Directory, you must register them with Oracle Internet Directory. This gives the directories access to change log objects stored in Oracle Internet Directory.

### About Directory Registration

To register a directory, you make an entry for it in Oracle Internet Directory. This entry is called a change subscription object, and it is placed under the following container in the Oracle Internet Directory schema:

```
cn=Subscriber Profile,cn=ChangeLog Subscriber,cn=Oracle Internet Directory
```

This change subscription object provides a unique credential for a directory to bind with Oracle Internet Directory and to retrieve changes from it.

Associate the change subscription object with the auxiliary object class `orclChangeSubscriber`, which has several attributes, two of them mandatory. The two mandatory attributes are:

| | |
|---|---|
| `userPassword` | Password to be used by the directory when accessing the change log object in Oracle Internet Directory |
| `orclLastAppliedChangeNumber` | Number of the change applied during the last synchronization. This attribute allows the directory to retrieve only the changes in Oracle Internet Directory it has not already applied. |

### Registering a Directory

To register a directory, use ldapadd. The following example uses an input file, named add.ldif, to create a change subscription object, `my_other_directory`, under the container `cn=Subscriber Profile,cn=ChangeLog Subscriber,cn=`Oracle Internet Directory.

- Edit file `add.ldif`:

```
dn: cn=my_other_directory,cn=Subscriber Profile,cn=ChangeLog Subscriber,
    cn=Oracle Internet Directory
userpassword:my_secret_code
orclLastAppliedChangeNumber: current_change_number_in_directory_before_
                             initial_boot_strapping
objectclass: orclChangeSubscriber
objectclass: top
```

- Add the entry:

```
ldapadd -h <host > -p < port > -f add.ldif
```

### Deregistering a Directory

To deregister a directory, use ldapdelete. Enter the following command:

```
ldapdelete -h host -p port cn=directory_name,cn=Subscriber Profile,
cn=ChangeLog Subscriber,cn=Oracle Internet Directory
```

## Task 3: Grant Directories Access to the Oracle Internet Directory Change Log Object Store

Once you have registered a directory with Oracle Internet Directory, you must grant it read access to the cn=changeLog entry in Oracle Internet Directory.

> **See Also:** Chapter 9, "Managing Directory Access Control" for instructions on setting access control policies

# 12

# Managing National Language Support (NLS)

Oracle Internet Directory National Language Support (NLS) enables you to store, process and retrieve data in native languages. It ensures that Oracle Internet Directory utilities and error messages automatically adapt to the native language and locale.

This chapter discusses NLS as used by Oracle Internet Directory and tells you the required NLS_LANG environment variables for the various components and tools in an Oracle Internet Directory environment.

> **See Also:** "National Language Support" on page 2-18 prior to configuring NLS

This chapter contains these topics:

# The NLS_LANG Environment Variable

The NLS_LANG parameter has three components—`language`, `territory`, and `charset`—in the form:

`NLS_LANG = language_territory.charset`

Each component controls the operation of a subset of NLS features.

| Component | Description |
|---|---|
| *language* | Specifies conventions such as the language used for Oracle messages, day names, and month names. Each supported language has a unique name—for example, American English, French, or German. The language argument specifies default values for the territory and character set arguments, so either (or both) `territory` or `charset` can be omitted.<br><br>If language is not specified, the value defaults to American English.<br><br>**See Also:** *Oracle8i National Language Support Guide*. for a complete list of languages |
| *territory* | Specifies conventions such as the default calendar, collation, date, monetary, and numeric formats. Each supported territory has a unique name; for example, America, France, or Canada.<br><br>If territory is not specified, the value defaults to America.<br><br>**See Also:** *Oracle8i National Language Support Guide*. for a complete list of territories |
| *charset* | Specifies the character set used by the client application (normally that of the user's terminal). Each supported character set has a unique acronym, for example, US7ASCII, WE8ISO8859P1, WE8DEC, WE8EBCDIC500, or JA16EUC. Each language has a default character set associated with it. Default values for the languages available on your system are listed in your operating system installation guide or administrator's guide.<br><br>Oracle Internet Directory requires all data to be stored in UTF-8.<br><br>**See Also:** *Oracle8i National Language Support Guide*. for a complete list of character sets |

**Note:** All components of the NLS_LANG definition are optional, that is, any item left out will default.

Also, if you specify `territory` or `charset`, you *must* include the preceding delimiter [underscore ( _ ) for `territory`, and period ( . ) for `charset`], otherwise the entire value will be parsed as a language name.

You can set NLS_LANG as an environment variable at the command line. The following are examples of legal values for NLS_LANG:

- `AMERICAN_AMERICA.UTF8`
- `JAPANESE_JAPAN.UTF8`

# Using NLS with LDIF Files

**See Also:** "LDAP Data Interchange Format (LDIF) Syntax" on page A-2

Attribute types are always ASCII strings that cannot contain multibyte characters. Oracle Internet Directory does not support multibyte characters in attribute type names. However, Oracle Internet Directory does support attribute *values* containing multibyte characters such as those in the simplified Chinese (.ZHS16GBK) character set.

Attribute values can be encoded in different ways to allow Oracle Internet Directory tools to interpret them properly. There are two scenarios:

- An LDIF file Containing Only ASCII Strings
- An LDIF file Containing UTF-8 Encoded Strings

## An LDIF file Containing Only ASCII Strings

In this scenario, character strings for attribute values are also in ASCII.

Because all tools use the UTF-8 character set by default, and ASCII is a proper subset of UTF-8, all tools can interpret these files. The same is true of keyboard input of values that are simply ASCII strings.

## An LDIF file Containing UTF-8 Encoded Strings

In this scenario, character strings for attribute values are also in UTF-8.

Because all tools use the UTF-8 character set by default, all tools can interpret these files. The same is true of keyboard input of values which are UTF-8 strings.

In such a file, some characters may be multibyte. Multibyte characters strings can be present in the LDIF files as attribute values or given as keyboard input. They can be encoded in their native character set or in UTF-8. They can also be BASE64 encoded representations of either the native or the UTF-8 string.

Consider the following cases:

- CASE 1: Native Strings (Non-UTF-8)
- CASE 2: UTF-8 Strings
- CASE 3: BASE64 Encoded UTF-8 Strings
- CASE 4: BASE64 Encoded Native Strings

Because the LDAP server understands and expects only UTF-8 encoded strings, cases 1, 3, and 4 need to undergo conversion to UTF-8 strings before they can be sent to the LDAP server.

### CASE 1: Native Strings (Non-UTF-8)

Use the `-E` argument in the command line tools, ldifwrite, and bulkmodify. Use the -encode argument in the bulkload and bulkdelete tools.

This example converts simplified Chinese native strings to UTF-8. The baseDN can be a simplified Chinese string:

```
ldapsearch -h my_host -p 389 -E ".ZHS16GBK" -b base_DN  -s base  'objectclass=*'
```

### CASE 2: UTF-8 Strings

No conversion is required.

### CASE 3: BASE64 Encoded UTF-8 Strings

You need to use neither the `-E` argument in the command line tools, ldifwrite, and bulkmodify, nor the `-encode` argument in bulkload and bulkdelete. Oracle Internet Directory tools automatically decode BASE64 encoded UTF-8 strings to UTF-8 strings.

### CASE 4: BASE64 Encoded Native Strings

Use the `-E` argument in the command line tools, ldifwrite, and bulkmodify. Use the `-encode` argument in the bulkload and bulkdelete tools.

Oracle Internet Directory tools automatically decode BASE64 encoded native strings to simple native strings. The native strings are then converted to the equivalent UTF-8 strings.

**Note:** In any given input file, only one language set may be used.

## Using NLS with Command Line Tools

The Oracle Internet Directory command line tools read keyboard input or LDIF file input in the following ways:

- ASCII characters only
- Non-ASCII input (native language character set)
- BASE64 encoded values of UTF-8 or native strings (from LDIF file only)

If the character set being given as input from an LDIF file or keyboard is not UTF-8, the command line tools need to convert the input into UTF-8 format before sending it to the LDAP server.

You enable the command line tools to convert the input into UTF-8 by specifying the `-E` argument when using each tool.

This section contains these topics:



### Specifying the -E Argument When Using Each Tool

The client tools always assume UTF-8 to be the character set unless otherwise specified by the `-E` argument. The BASE64-encoded values are decoded, and then the decoded buffer is converted to UTF-8 if the `-E` argument is specified. For example, if you specify `-E ".ZHS16GBK"`, then the decoded buffer is converted from simplified Chinese to UTF-8 before being sent to the LDAP server.

Specifying the `-E` argument ensures that proper character set conversion can occur from the character set you specify for the `-E` argument (`-E ".`*`character_set`*`"`) to the.UTF-8 character set.

The command line tools use the -E argument to process the input in the character set specified for the -E argument. They display their output in the character set specified in the NLS_LANG environment variable.

For example, to add entries from an LDIF file encoded in the simplified Chinese character set (.ZHS16GBK) by using ldapadd, type:

```
ldapadd -h myhost -p 389 -E ".ZHS16GBK" -f my_ldif_file
```

In this example, the ldapadd tool converts the characters from ".ZHS16GBK" (simplified Chinese character set) to ".UTF8" (UTF-8 character set) before they are sent across the wire to the LDAP server.

## Examples: Using the -E Argument with Command Line Tools

The following table provides additional examples of how to use the -E argument correctly for each command line tool. In each example, the command converts data from simplified Chinese, as specified by the value ".ZHS16GBK", to UTF-8. For example, in each command, the values for the -D and -w options are in simplified Chinese. Specifying the -E argument converts them to UTF-8.

Note that, in the examples in the following table, we do not show any actual characters belonging to .ZHS16GBK character set. These examples would, therefore, work without the -E argument. However, if the argument values contained actual characters in the .ZHS16GBK character set, then we would need to use the -E argument.

> **See Also:** Appendix A, "Syntax for LDIF and Command Line Tools" for syntax and usage notes for each of the command line tools

| Tool | Example |
|---|---|
| ldapbind | `ldapbind -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |
| ldapsearch | `ldapsearch -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |
| ldapadd | `ldapadd -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |
| ldapaddmt | `ldapaddmt -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |

| Tool | Example |
|---|---|
| ldapmodify | `ldapmodify -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |
| ldapmodifymt | `ldapmodifymt -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |
| ldapdelete | `ldapdelete -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password` |
| ldapcompare | `ldapcompare -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password -b ou=Construction,ou=Manufacturing,o=acme,c=us -a title -v manager` |
| ldapmoddn | `ldapmoddn -h my_host -p 389 -E ".ZHS16GBK" -D o=acme,c=us -w my_password -b "cn=Franklin Badlwins,ou=Construction,ou=Manufacturing,c=us,o=acme" -N ou=Contracting,ou=Manufacturing,o=acme,c=us -r` |

## Setting NLS_LANG in the Client Environment

If the output required by the client is UTF-8, then you do not need to set the NLS_LANG environment variable. In this case, the NLS_LANG environment variable defaults to `.UTF8`, and both the input path from client to server, and the output path from server to client, do not require any character set conversion.

If the output required by the client is *not* UTF-8, then you must set the NLS_LANG environment variable. This ensures that proper character set conversion can occur from the UTF-8 character set to the character set required by the client.

For example, if the NLS_LANG environment variable is set to the simplified Chinese character set, then the command line tool displays output in that character set. Otherwise the output defaults to the UTF-8 character set.

**Note:** If you are using Windows NT, then, to use the command line tools after server startup, you must reset NLS_LANG in an MS-DOS window. Set it to the character set that matches the code page of your MS-DOS session. (UTF-8 cannot be used.) See the *Oracle8i Installation Guide for Windows NT* for more information on which character set to use for command line tools in an MS-DOS session.

If you are using a pre-installed Oracle8*i* Release 8.1.7 database with Oracle Internet Directory, then you must also set the database character set to UTF-8. See the *Oracle8i National Language Support Guide* and *Oracle8i Installation Guide for Windows NT* for Windows NT for more information.

Be careful not to change the NLS_LANG parameter value in the registry.

# Using NLS with Bulk Tools

Oracle Internet Directory ensures that the reading and writing of text data from and to LDIF files are done in UTF-8 encoding as specified by the LDAP standard.

This section provides an example of the argument you use for each of the following bulk tools:

- Using NLS with bulkload
- Using NLS with ldifwrite
- Using NLS with bulkdelete
- Using NLS with bulkmodify

**See Also:** "Bulk Tools Syntax" for a list of arguments for each bulk tool

## Using NLS with bulkload

Add to the command the argument `-encode "`*`character_set`*`"` where the input LDIF file is encoded in *`"character_set"`*.

For example:

```
bulkload.sh -connect net_service_name -encode ".ZHS16GBK" my_ldif_file
```

## Using NLS with ldifwrite

The ldifwrite utility always writes BASE64 encoded values for multibyte strings.

The BASE64 encoding could be of the UTF-8 strings as they are stored in the directory server, or of native strings as specified by the NLS_LANG environment variable setting when running ldifwrite.

For example:

```
ldifwrite -c net_service_name -b baseDN -f output_file
```

In this example, if the NLS_LANG environment variable is not set, or is set to *language_territory*.UTF8, then the output LDIF file will contain BASE64-encoded UTF-8 strings for any multibyte characters.

To reload this LDIF file into the directory by using ldapaddmt, use the following syntax:

```
ldapaddmt -h host -p port -f output_file
```

In the above case, the -E argument is not required because the decoded BASE64 strings are already UTF-8-encoded and can be readily sent to the server.

If the NLS_LANG environment variable is set to a character set other than UTF-8—for example, ".ZHS16GBK"—then the output LDIF file will contain a BASE64 encoded value of simplified Chinese (.ZHS16GBK) strings.

To reload this LDIF file into the directory using ldapaddmt, use the following syntax:

```
ldapaddmt -h host -p port -E ".ZHS16GBK" -f my_input_file.LDIF
```

In the above case the -E argument is required because the decoded BASE64 strings are simplified Chinese, which need to be converted to UTF-8 strings before being sent to the server.

## Using NLS with bulkdelete

Add -encode ".*character_set*" to the command the argument.

For example:

```
bulkdelete.sh -connect net_service_name -encode ".ZHS16GBK" -base
"ou=manufacturing,o=acme,c=us"
```

In this case the value for the `-base` option could be in the `ZHS16GBK` native character set, that is, simplified Chinese.

## Using NLS with bulkmodify

Add `-E ".character_set"` to the command the argument.

For example:

```
bulkmodify.sh -c net_service_name -E ".ZHS16GBK" -b ou=manufacturing,o=acme,c=us
-r title -v Foreman -f filter
```

In this example, values for the `-b`, `-v`, and `-f arguments` can be specified using the simplified Chinese character set.

# Part III

## Deploying Oracle Internet Directory

Part III discusses deployment considerations. It contains these chapters:

# 13

# Deployment Considerations

This chapter discusses issues to consider when deploying Oracle Internet Directory. It helps you assess enterprise directory requirements and make effective deployment choices. Although the recommendations in this chapter are primarily for directories in medium to large enterprises and Internet Service Providers (ISPs), the principles apply to other environments as well.

This chapter contains these topics:

# The Expanding Role of Directories

Today, most enterprises are at various stages of deploying centralized and consolidated LDAP-compliant directories. Some have had non-LDAP-compliant directories—for example, NDS or ISO X.500—and are now converting to the corresponding LDAP-enabled versions. This is either to accommodate LDAP-reliant Internet clients, such as those embedded in Web browsers, or to consolidate the increasing number of platforms and services that use directories.

The increased numbers of LDAP-enabled applications make availability and performance requirements for LDAP-compliant directories critical. Most environments need to update their deployments.

Enterprises should plan a robust and flexible deployment to accommodate:

- The increased volume of information in the directory
- The number of applications that rely on the directory
- Such load characteristics as concurrent access and throughput

Choosing a directory product that can function as an enterprise backbone is a serious decision. As the directory becomes more central to the operation of the network and its services, deployment choices also become critical.

# Logical Organization Of Directory Information

Establishing an effective policy for **directory information tree (DIT)** structure and naming requires enterprise-wide coordination and planning. For example, the following questions can arise:

- How do you choose your enterprise directory naming and organization?
- Should the choice reflect the corporate organizational structure or geographic and national boundaries?
- Does the choice work seamlessly for NOS directories such as Novell's eDirectory solution and Microsoft Active Directory?

This section contains these topics:

## Directory Entry Naming

Typically, most enterprises have a Human Resources department that establishes rules for assigning unique names and numbers for employees. When choosing a unique naming component for directory entries, it is good to exploit this administrative infrastructure and use its policies. By contrast, the advantage of making DNs more "user friendly" is outweighed by the proliferation of administrative policies it would require.

## DIT Hierarchy and Structure

A DIT is hierarchical in structure, similar to the DNS (Domain Name System). It is possible to organize the DIT to reflect any logical hierarchy associated with an enterprise. The choice should accommodate the following:

- The DIT structure and naming policies for the enterprise as a whole should be compatible with the rules and restrictions of departmental NOS directories. For example, some directory products define domains, and then require organizational units and localities to be logically subordinate to those domains. Also, some directory products require directory name uniqueness within a domain, even for entries that are not **siblings**.
- The directory organization should facilitate clear and effective access control and replication policies. In an enterprise where delegation of **ACL** administration is required, it is better to organize the DIT to reflect the data ownership boundaries.

  For example, consider a corporation which has an autonomous data center for each major geographic region: one for the Americas (North and South), one for Europe, and one for Asia Pacific. Suppose that this corporation wants to consolidate its global directory, while retaining the administrative autonomy of its regional data centers. It should organize the directory in **naming contexts** corresponding to each region. This makes it easier to develop access control and replication policies that suit regional needs.
- It may be tempting to organize the directory hierarchy to reflect either the corporate divisional structure or the organizational hierarchy. Usually, this is not advisable because most corporations undergo frequent reorganization and divisional restructuring. It is more manageable to capture a person's organizational information as an attribute of the person's directory entry.

# Physical Distribution: Partitions and Replicas

You can distribute directory data in two ways:

- By maintaining the entire directory on one server
- By hosting different naming contexts on different servers and maintaining **knowledge references**, also called referrals, between them

**See Also:** "Distributed Directories: An Overview" on page 2-26

This section contains these topics:



## An Ideal Deployment

In an ideal world, it would be simpler and more secure to store all naming contexts in a central consolidated directory server. The problem is that this central directory server would then be a single point of failure.

A simple solution might be to implement redundant LDAP servers and their associated databases. However, even redundancy might not provide the needed connectivity, accessibility, and performance that most global organizations need at all their regions and sites. These requirements might, in fact, call for replicas physically located at various regions across the corporate geography.

If Oracle Internet Directory supported only single-master configuration, then logical consolidation of the directory would be difficult. Each region or group would want to store the master replica for the naming context on which that group relies. This, in turn, could mean a lack of uniformity in the administrative policies among the partitions. Administrators would need to use a different data management procedure for each partition.

However, Oracle Internet Directory's multimaster replication makes logical consolidation of the directory easier. It allows "update anywhere" configurations, which makes consolidating the directory more efficient and less costly than maintaining multiple partitions.

Here is a simple and practical recommendation for a robust centralized corporate directory:

- Establish a network of two or more directory nodes, each holding all the naming contexts. Set up these nodes in a multimaster configuration.
- Deploy these individual nodes, one in each geographic region, to suit the corporate data network connectivity. For example, if a region is connected to the rest of the network by way of a slow link, then it is better to locate a dedicated directory server for use by the clients in that region.
- Individually configure each regional server for failover and recovery.

Remember: Even if all the naming contexts are consolidated, you can still achieve administrative autonomy for various logical naming contexts. You do this by establishing appropriate access control policies at the root of each naming context.

**See Also:** "Failover Considerations" on page 13-7 for a discussion of redundancy

## Partitioning Considerations

A directory with too many **partitions** generally has more administrative overhead than benefits. This is because each partition requires you to plan backup, recovery, and other data management functions.

Typically, the reasons for maintaining partitions are:

- They correspond to administrative and data ownership boundaries that are better left independent
- The enterprise network has regions that are connected with expensive or low-speed links and many partitions have only local access needs
- The lack of availability of a partition does not have a larger impact
- Maintaining an entire corporate directory in a certain region is too expensive

When you use partitioning, interconnect the partitions with **knowledge references**.

---

**Note:** LDAP does not support automatic chaining of knowledge references by the LDAP server. The majority of client side LDAP APIs support client-driven knowledge reference chasing. However, there is no guarantee that knowledge references will be supported in all the LDAP tools. The lack of consistent knowledge reference support across all available tools is a factor to consider before deciding to use partitions.

---

## Replication Considerations

LDAP directory replication architecture is based on a loose consistency model: Two replicated nodes in a **replication agreement** are not guaranteed to be consistent in real time. This increases the overall flexibility and availability of the directory network, because a client can modify data without all interconnected nodes being available. Suppose, for example, that one node is unavailable or heavily loaded. With multimaster replication, the operation can be performed on an alternate node, and all interconnected nodes synchronize in due course.

There are many reasons to implement a replicated network, including the following:

- Local accessibility and performance requirements

  Most corporations have operations in many regions in the world, and those operations need a common directory. Suppose that the regions were interconnected with low bandwidth links involving multiple intermediate routers. A client accessing a directory server from outside the region could experience a very high **latency**, and even inadequate **throughput**.

  In such cases, a regional replica—enabled by multimaster replication to receive updates— is essential. Moreover, the replication data transfer can be scheduled for off-peak hours in the underlying **Advanced Symmetric Replication (ASR)**.

- Load balancing

  When directory access exceeds the capacity of an existing server, an additional server must share the load. With Oracle Internet Directory, two such systems can be deployed in a multimaster replication mode. In fact, even when planning the directory deployment to meet a specific estimated load, it can be less costly to maintain two relatively low-end systems than one high-end system. In addition to load balancing, such configurations also contribute to higher system availability.

- Failure tolerance and higher overall system availability

  One of the most important reasons to implement directory replication is to increase overall system availability. When one server is unavailable, the traffic can be routed to other available servers. This can be transparent to clients.

## Failover Considerations

Because a directory service has a critical function in an enterprise, deployment should take failure recovery and high availability into consideration. This includes developing backup and recovery strategies for individual nodes.

In addition to multimaster replication, consider the following failover and high-availability options for potential deployment at any Oracle Internet Directory installation:

- Intelligent Client Failover

  All LDAP clients connecting to Oracle Internet Directory can maintain a list of alternate server instances of Oracle Internet Directory to contact if their connection with a given server instance is abruptly broken.

- Intelligent Network Level Failover

  There are several hardware and software solutions that can detect the failure of the system hosting Oracle Internet Directory. These solutions can intelligently reroute future connection requests to an alternate server. Some of these solutions balance the load of incoming connection requests with alternate servers, while also providing the necessary failover capabilities.

Because Oracle Internet Directory is a client of Oracle8*i*, other failover technologies, such as Oracle Parallel Server, are also available.

> **See Also:** Chapter 16, "High Availability And Failover" for further details about high-availability and failover options available with Oracle Internet Directory.

## About Capacity Planning, Sizing, and Tuning

When estimating enterprise-wide and regional requirements for directory usage, plan for future needs. Depending on other configuration choices for replication and failover, there could be more than one directory node, each with its own load and capacity requirements. In this case, you must individually size each directory node.

As an enterprise increases its directory usage, more applications rely on Oracle Internet Directory to serve their requests in a timely manner. Ensure that the Oracle Internet Directory installation can live up to the performance and capacity expectations of those applications.

You can influence the capacity and performance of a given Oracle Internet Directory installation in two phases of the deployment process:

- Planning phase

  During this phase, gather the requirements of all directory users and establish a unified performance and capacity requirement. This consists of capacity planning and system sizing.

- Implementation phase

  Once you have the hardware, tune the Oracle Internet Directory software stack for best use of the hardware resources. This improves the performance of Oracle Internet Directory and of the LDAP client applications.

This section contains these topics:



## Capacity Planning

Capacity planning is the process of determining performance and capacity requirements. You base these on typical models of directory usage in the enterprise.

When trying to estimate the required capacity of an Oracle Internet Directory installation, consider:

- The type of LDAP client applications
- The number of users accessing those applications
- The nature of LDAP operations those applications perform
- The number of entries in the DIT
- The type of operations performed against the Oracle directory server
- The number of concurrent connections to the Oracle directory server
- The peak rate at which operations need to be performed by the Oracle directory server
- The average latency of operations required under peak load conditions

While estimating these details, allow room for future increases in directory usage.

## Sizing Considerations

Once you have established the fundamental capacity and performance requirements, translate them into system requirements. This is called system sizing. Some of the details to consider in this phase are:

- The type and number of CPUs for the Oracle Internet Directory server computer
- The type and size of disk subsystems for the Oracle Internet Directory server computer
- The amount of memory required for the Oracle Internet Directory server computer
- The type of network used for LDAP messages from the clients

Based on current experience, the following table indicates the approximate level of CPU power required for various deployment scenarios for Oracle Internet Directory:

| Usage | Num CPUs | SPECint_rate95 baseline | System |
|---|---|---|---|
| Departmental | 2 | 60 to 200 | Compaq AlphaServer 8400 5/300 (300Mhz x 2) |
| Organization wide | 4 | 200 to 350 | IBM RS/6000 J50 (200MHz x 4) |
| Enterprise wide | 4+ | 350+ | Sun Ultra 450 (296 MHz x 4) |

The amount of disk space required for an installation of Oracle Internet Directory is directly proportional to the number of entries stored in the DIT. The following table gives the approximate disk space requirements for variously sized DITs.

| Number of Entries in DIT | Disk Requirements |
|---|---|
| 100,000 | 450MB to 650MB |
| 200,000 | 850MB to 1.5GB |
| 500,000 | 2.5GB to 3.5GB |
| 1,000,000 | 4.5GB to 6.5GB |
| 1,500,000 | 6.5GB to 10GB |
| 2,000,000 | 9GB to 13GB |

The data in this table makes the following assumptions:

- There are approximately 20 cataloged attributes
- There are approximately 25 attributes for each entry
- The average size of an attribute is approximately 30 bytes

The amount of memory required for Oracle Internet Directory is mostly governed by the amount of database cache that a deployment site desires. Often, the size of the database cache is directly proportional to the number of entries in the DIT. The following table provides estimates of the memory requirements for various DIT sizes:

| Directory Type | Number of Entries | Minimum Memory |
|---|---|---|
| Small | Less than 600,000 | 512MB |
| Medium | 600,000 to 2,000,000 | 1GB |
| Large | Greater than 2,000,000 | 2GB |

**See Also:** Chapter 14, "Capacity Planning."

## Tuning Considerations

Oracle Corporation recommends that you properly tune Oracle Internet Directory before using it in a production environment. Before tuning, ensure that there are adequate testing mechanisms and sample data in the directory to simulate a real world usage scenario. Perhaps you can use the applications that rely on the directory for testing purposes.

Any tool for testing the performance of Oracle Internet Directory must be able to show:

- The overall throughput it is noticing
- The average latency of operations

In this way, the tool provides a feedback mechanism for determining the effects of tuning and providing direction to the overall tuning effort.

Some of the commonly tuned properties of an Oracle Internet Directory installation include:

- CPU usage

  This is determined, to a large extent, by:

  – The number of Oracle directory servers

  – The number of database connections opened by each server

  On the one hand, too large a number of Oracle directory servers and database connections can cause too much contention for available CPU resources. On the other hand, too small a number of Oracle directory servers and database connections can leave much of the CPU power under-utilized. Consider adjusting these numbers to the appropriate levels based on available CPU resources and the expected peak load.

- Memory usage

  The main consumer of memory in an Oracle Internet Directory installation is the database cache, which is part of the **SGA**. In some cases, allocating a very large database cache can eliminate much disk I/O for Oracle data files. However, it can also cause paging, which is detrimental to performance. Alternatively, having a small database cache causes too much disk I/O, and that is also detrimental to performance. Tune the memory usage of the system so that all consumers of memory in the system can get physical memory without needing to use paging.

- Disk usage

  Because all of the data served by Oracle Internet Directory resides in database tablespaces, pay attention to any tuning that can increase the I/O throughput. Common techniques for disk tuning include:

  – Balancing tablespaces on different logical and physical drives

  – Striping logical volumes onto multiple physical volumes

  – Distributing disk volumes across multiple I/O controllers

  **See Also:** Chapter 15, "Tuning" for further details on various tuning tips and techniques

# 14

# Capacity Planning

Capacity planning is the process of assessing applications' directory access requirements and ensuring that the Oracle Internet Directory has adequate computer resources to service requests at an acceptable rate. This chapter explains what you need to consider when doing capacity planning. It guides you through an example of a directory deployment for an email messaging application in a hypothetical company called Acme Corporation

This chapter contains these topics:

# About Capacity Planning

If Oracle Internet Directory and the corresponding Oracle8*i* database are running on the same computer, then these are the configurable resources that capacity planners need to consider:

- I/O subsystem (the type and size)
- Memory
- Network connectivity
- CPUs (speed and quantity)

When you plan to acquire hardware for Oracle Internet Directory, you should ensure that all components—such as CPU, memory, and I/O—are effectively used. Generally, good memory usage and a robust I/O subsystem are sufficient to keep the CPU busy.

Any new installation of the Oracle Internet Directory needs two things to be successful:

- Adequate hardware resources so that the installed system can satisfy user demands at peak load rates
- A well tuned system—hardware and software—that makes the best use of available resources, one that squeezes the maximum performance out of available hardware

We begin by looking at an example of a directory deployment for an email messaging application in a hypothetical company called Acme Corporation. As we examine each component of the capacity plan, we will apply our recommendations to the example of Acme Corporation.

### Terms Used Throughout This Chapter

| | |
|---|---|
| Throughput | The overall rate at which directory operations are being completed by Oracle Internet Directory. This is typically represented as "operations per second." |
| Latency | The time a client has to wait for a given directory operation to complete |
| Concurrent clients | The total number of clients that have established a session with Oracle Internet Directory |
| Concurrent operations | The amount of concurrent operations that are being executed on the directory from all of the concurrent clients. Note that this is not necessarily the same as the concurrent clients because some of the clients may be keeping their sessions idle. |

## Getting to Know Directory Usage Patterns: A Case Study

The ability to assess the potential load on Oracle Internet Directory is very important for developing an accurate capacity plan. Let us examine the email messaging software employed by our hypothetical company, Acme Corporation. The email messaging software in this example is based on Internet Message Access Protocol (IMAP). There are two main types of software that access Oracle Internet Directory:

- The IMAP clients, which will validate email addresses within the company before sending the mail to the IMAP server. These clients include software programs like Netscape Messenger and Microsoft Outlook.
- The messaging software itself, also called the Mail Transfer Agent (MTA), which will look up the directory to route mail from the outside world to internal mailboxes as well as route internal mails to company-wide distribution lists.

Let us assume that the private aliases and private distribution lists of individual users are also stored in the directory. Let us further make the following assumptions, which will allow us to guess the size of the directory:

| | |
|---|---|
| Total user population | 40,000 |
| Average number of private aliases per person | 10 |
| Average number of private distribution lists per person | 10 |
| Total number of public distribution lists | 4000 |
| Total number of public aliases in the company | 1000 |
| Number of attributes in each entry in the directory related to this application | 20 |
| Number of cataloged attributes | 10 |

Based on the above assumptions, we can derive the overall count of entries in Oracle Internet Directory as:

| | |
|---|---|
| User entries | 40,000 (these represent the users themselves) |
| Private aliases of users | 40,000 x 10 = 400,000 entries |
| Private distribution lists of users | 40,000 x 10 = 400,000 entries |
| Company wide distribution lists | 4000 |
| Company wide aliases | 1000 |

The above assumptions will yield a directory population of about one million entries. Given the user population and the directory population, let us then analyze usage patterns so that we can derive performance requirements from them. A typical user tends to send an average of 10 emails per day and receives an average of 10 emails a day from the outside world. Assuming that there are, on an average, five recipients for each email being sent by a user, this would result in five directory lookups for each email.

The following table summarizes all the possible directory lookups that can happen in one day:

| Type of Directory Lookup | Number of Directory Lookups In One Day |
|---|---|
| The Mail Transfer Agent (MTA) processing outbound mail from each user | 5x10x40,000 = 2,000,000 |
| The MTA processing mails from the outside world | 10x40,000 = 400,000 |
| All other directory lookups (like IMAP clients validating certain addresses etc.) | 800,000 |

Summing up, the total number of directory lookups per day would be about 3,200,000 (3.2 million) directory lookups per day. If these directory lookups were spread out uniformly along the day, it would require about 37 directory lookups per second (133,333 lookups per hour). Unfortunately, we will never have this case. Usage analysis of the current email system over a period of 24 hours shows the pattern illustrated in Figure 14–1.

***Figure 14–1 Usage Analysis of Current Email System***

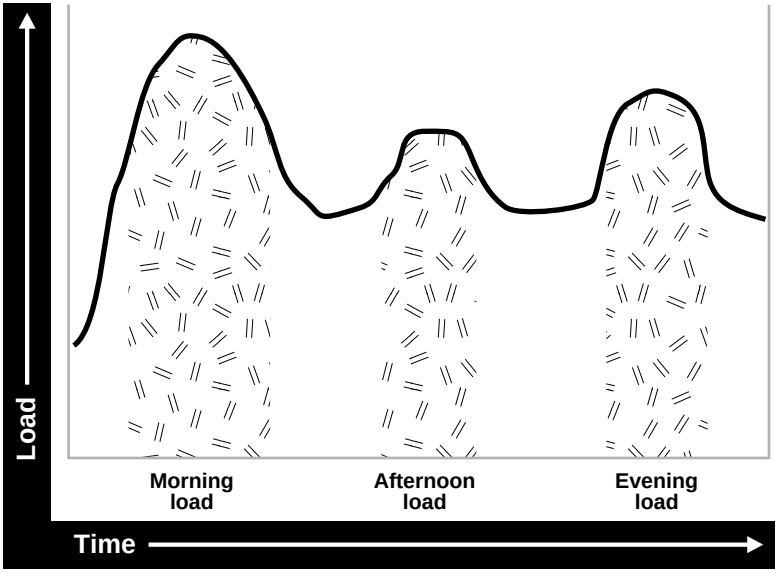

The email system (and Oracle Internet Directory) is stressed at its peak in the mornings. There are other usage peaks as well—one close to lunch time, and one near the end of business day. However, it is in the mornings that the Oracle Internet Directory is stressed the most.

Let us assume that 90 percent of all the directory lookups happen during normal working hours. Let us now split up the working hour load into the following categories (assuming an 8 hour workday):

| | |
|---|---|
| Morning load | 65%: 0.90 x 0.65 x 3,200,000 = 1,872,000 lookups for 2 hours (936,000 lookups per hour) |
| Afternoon load | 10%: 0.90 x 0.10 x 3,200,000 = 288,000 lookups for 1 hour (288,000 lookups per hour) |
| Evening load | 20%: 0.90 x 0.20 x 3,200,000 = 576,000 lookups for 2 hours (288,000 lookups per hour) |

The above calculations indicate that the Oracle Internet Directory in this case should be designed to handle the peak load of 936,000 lookups per hour.

Now that we know the data-set size as well as the performance requirements, we can now look into individual components of the installation and estimate good values for each.

# I/O Subsystem Requirements

This section contains these topics:

- About the I/O Subsystem
- Rough Estimates of Disk Space Requirements
- Detailed Calculations of Disk Space Requirements

## About the I/O Subsystem

The I/O subsystem can be compared to a pump that pumps data to the CPUs to enable them to execute workloads. The I/O subsystem is also responsible for data storage. The main components of an I/O subsystem are arrays of disk drives controlled by disk controllers.

It is important to consider performance requirements when you size the I/O subsystem, rather than size based only on storage requirements. Although disk drives have increased in size, the throughput—that is, the rate at which the disk

drive pumps data—has not increased in proportion. In sizing calculations for the I/O subsystem, you should use the following factors as input:

- The size of the database
- The number of CPUs on the system
- An initial estimation of the workload on the Oracle Internet Directory
- The rate at which the disk can pump data
- Space needed to stage data prior to load
- Space needed for index creation and sort activities

Given a range of I/O subsystems, you should always opt for the highest throughput drives. Typically, one can maximize the I/O throughput by one or more of the following techniques:

- Striping logical volumes so that the I/O operations use multiple disk spindles
- Putting different tablespaces in different logical and physical disk volumes
- Distributing the disk volumes on multiple I/O controllers

Some guidelines for organizing Oracle Internet Directory-specific data files are provided in Chapter 15, "Tuning". Depending on the tolerance of disk failures, different levels of Redundant Arrays of Inexpensive Disks (RAID) can also be considered.

Assuming that the decision has been made to get the best possible I/O subsystem, we focus the next section on deriving sizing estimates for the disks themselves.

## Rough Estimates of Disk Space Requirements

You can use the following table to derive a rough estimate of the overall disk requirement:

| Number of Entries in DIT | Disk Requirements |
|---|---|
| 100,000 | 450MB to 650MB |
| 200,000 | 850MB to 1.5GB |
| 500,000 | 2.5GB to 3.5GB |
| 1,000,000 | 4.5GB to 6.5GB |
| 1,500,000 | 6.5GB to 10GB |
| 2,000,000 | 9GB to 13GB |

The data shown in the previous table makes the following assumptions:

- There are about 20 cataloged attributes.
- There are about 25 attributes per entry.
- The average size of an attribute is about 30 bytes.

Going back to our example of Acme Corporation, since our directory population is about one million, this would imply that our disk requirements are approximately 4.5 GB to 6.5 GB. Note that the assumptions made for Acme Corporation regarding the number of cataloged attributes are different, but the previous table should give an approximate figure of the size requirements.

Since the directory may be deployed for a wide variety of applications, these assumptions need not necessarily hold true for all possible situations: There might be cases where the size of attributes is large, the number of attributes per entry is large, extensive use of ACIs has been made, or the number of cataloged attributes is very high. For such cases, we present simple arithmetic procedures in the following section which will allow the planners to get a more detailed perspective of their disk requirements.

## Detailed Calculations of Disk Space Requirements

Because Oracle Internet Directory stores all of its data in an Oracle8*i* database, the sizing for disk space is primarily a sizing of the underlying database. Oracle Internet Directory stores its data in the following tablespaces:

| | |
|---|---|
| OLTS_ATTR_STORE | Stores all of the attributes for all entries in the DIT |
| OLTS_IND_ATTRSTORE | Stores the indices pertaining to attributes in the directory |
| OLTS_CT_DN | Stores the distinguished name catalog |
| OLTS_IND_CT_DN | Stores the indices pertaining to the DN catalog |
| OLTS_CT_CN | Stores the common name catalog |
| OLTS_CT_OBJCL | Stores the ObjectClass catalog |
| OLTS_CT_STORE | Stores all the remaining (including user-defined) catalogs |
| OLTS_IND_CT_STORE | Stores the indices pertaining to the user-defined catalogs |
| OLTS_DEFAULT | Stores all of the data pertaining to the administration of the Oracle Internet Directory as well as the data used for replication support |
| OLTS_TEMP | Used for creating various indices on the tables. It should be large enough so that all index creations can go through. |
| SYSTEM | Required by Oracle8*i* database for various book-keeping purposes. Typically, its size remains constant at about 300MB. |

This section presents simple arithmetic procedures to determine the size requirements of each of the tablespaces shown above. All of the size calculations are based on the following variables:

| Variable Name | Description |
|---|---|
| *num_entries* | Total number of entries in the directory |
| *attrs_per_entry* | Average number of attributes per directory entry |
| *avg_attr_size* | Average size of the attribute in bytes |
| *avg_dn_size* | Average size of the DN of an attribute in bytes |
| *objectclass_per_entry* | Average number of object classes that an entry belongs to |
| *objectclass_size* | Average size of the name of each objectclass in bytes |
| *num_cataloged_attrs* | Number of cataloged attributes used in the entries |
| *entries_per_catalog* | Average number of entries per catalog table. This is required because not all cataloged attributes will be present in all entries in the DIT. |
| *change_log_capacity* | Number of changes that we wish to buffer for replication purposes |
| *num_acis* | Overall number of ACIs in the directory |
| *num_auditlog_entries* | Number of auditlog entries to store in the directory |
| *db_storage_ovhd* | Overhead of storing data in tables. This overhead corresponds to the relational constructs as well as operating system specific overhead. A value of 1.3 for this variable would represent a 30 percent overhead. The minimum value for this variable is 1. |
| *db_index_ovhd* | Overhead of storing data in indices. This overhead corresponds to the relational constructs as well as the operating system specific overhead. A value of 5 for this variable would represent a 400 percent overhead. The minimum value of this variable is 1. |
| *factor_of_safety* | Multiplier for accommodating growth and errors in calculations. A value of 1.3 for this variable would represent a 30 percent factor of safety. The minimum value for this variable is 1. |

Using the variables shown in the preceding table, the size of individual tablespaces can be calculated as follows:

| Tablespace Name | Size |
|---|---|
| OLTS_ATTR_STORE | *num_entries* * *attrs_per_entry* * *avg_attr_size* * *db_storage_ovhd* |
| OLTS_IND_ATTRSTORE | *num_entries* * *attrs_per_entry* * 30 |
| OLTS_CT_DN | *num_entries* * 2 * *avg_dn_size* |
| OLTS_IND_CT_DN | *num_entries* * 2 * *(avg_dn_size* + 30*)* |
| OLTS_CT_CN | *num_entries* * *avg_dn_size* * *db_storage_ovhd* |
| OLTS_CT_OBJCL | *(num_entries * objectclass_per_entry * objectclass_size *db_storage_ovhd) + (num_auditlog_entries* * 2 * *avg_dn_size * db_storage_ovhd)* |
| OLTS_CT_STORE | *(entries_per_catalog * num_cataloged_attrs * avg_attr_size * db_storage_ovhd) + (num_entries * objectclass_per_entry * objectclass_size * db_storage_ovhd)* |
| OLTS_IND_CT_STORE | *(entries_per_catalog * num_cataloged_attrs * avg_attr_size * db_index_ovhd) + (num_entries * objectclass_per_entry * objectclass_size * db_index_ovhd) + (num_acis* * 1.5 * *avg_dn_size* * *db_index_ovhd) + (num_auditlog_entries* * 2 * *avg_dn_size * db_index_ovhd)* |
| OLTS_DEFAULT | *(change_log_capacity* * 4 * *avg_attr_size * db_storage_ovhd * db_index_ovhd) + (num_entries* * 5*)* |
| OLTS_TEMP | *(*size of *OLTS_IND_ATTR_STORE) + (*size of *OLTS_IND_CT_STORE)* |
| SYSTEM | 300 MB |

Using the arithmetic operations shown in the preceding table, one can compute the exact space requirements for a wide variety of Oracle Internet Directory deployment scenarios. The sum of the sizes of each of the tablespaces should yield the overall database disk requirement. One can optionally multiply that by the "factor_of_safety" variable to get a figure that can compensate for unforeseen circumstances.

Going back to our example of Acme Corporation, we can assign values to each of the variables based on the requirements stated in previous sections. The following table illustrates the values of each variable introduced in this section for Acme Corporation.

| Variable Name | Value |
|---|---|
| *num_entries* | 1,000,000 |
| *attrs_per_entry* | 20 |
| *avg_attr_size* | 32 bytes |
| *avg_dn_size* | 40 bytes |
| *objectclass_per_entry* | 5 (each entry belongs to an average of 5 object classes) |
| *objectclass_size* | 10 bytes |
| *num_cataloged_attrs* | 10 |
| *entries_per_catalog* | 1,000,000 |
| *change_log_capacity* | 80,000 changes (2 per user) |
| *num_acis* | 80,000 ACIs (2 per user) |
| *num_auditlog_entries* | 1000 |
| *db_storage_ovhd* | 1.4 (40% overhead) |
| *db_index_ovhd* | 5.0 (400% overhead) |
| *factor_of_safety* | 1.5 (50% factor of safety) |

If we now plug these values into the equations described earlier, we get the following values:

| Tablespaces Name | Size in Bytes | Size in MB | Size in MB (with factor of safety) |
|---|---|---|---|
| OLTS_ATTRSTORE | 896000000 | 875 | 1313 |
| OLTS_IND_ATTRSTORE | 600000000 | 586 | 879 |
| OLTS_CT_DN | 80000000 | 78 | 117 |
| OLTS_IND_CT_DN | 140000000 | 137 | 205 |
| OLTS_CT_CN | 56000000 | 55 | 82 |
| OLTS_CT_OBJCL | 70112000 | 68 | 103 |
| OLTS_CT_STORE | 518000000 | 506 | 759 |
| OLTS_IND_CT_STORE | 1874400000 | 1830 | 2746 |
| OLTS_DEFAULT | 76680000 | 75 | 112 |
| OLTS_TEMP | 2474400000 | 2416 | 3625 |
| SYSTEM | 307200000 | 300 | 450 |
| **Total Size** | **7092792000** | **6927** | **10390** |

The table above shows that the estimated size of the database for Acme Corporation would be about 6.9 GB. With a 50 percent factor of safety, this would jump to 10.4GB. If all of the data is being loaded in bulk, then the bulkload tool of Oracle Internet Directory would require an additional 50 percent of space occupied by the database to store its temporary files. For Acme Corporation, this would add about 2.25 GB to the total space requirement.

## Memory Requirements

Memory is used for a number of distinct tasks by any database application, including Oracle Internet Directory. If memory resources are insufficient for any of these tasks, the bottleneck causes the CPUs to work at lower efficiency and system performance to drop. Furthermore, memory usage increases in proportion to the number of concurrent connections to the database and the number of concurrent users of the directory.

The memory available to processes comes from the virtual memory on the system, which is somewhat more than available physical memory. If the sum of all active

memory usage exceeds the available physical memory on the system, the operating system may need to store some of the memory pages on disk. This is called paging. Paging can degrade performance if memory is too oversubscribed. Generally, you should not exceed 20 percent over-subscription of physical memory. If paging occurs, you need either to scale back memory usage by processes or to add more physical memory. Keep in mind the trade-offs: There are physical limits to the amount of memory you can add, but scaling back on per-process memory usage can significantly degrade performance.

The main consumer of memory is the database buffer cache within the **System Global Area (SGA)**. The more memory allocated to this, the better will be the buffer cache hit ratio. A good buffer cache hit ratio will result in good database performance which in turn will result in good performance of the Oracle Internet Directory.

> **See Also:** Chapter 15, "Tuning" for further information on SGA tuning

The following table gives minimum memory requirements for different directory configurations:

| Directory Type | Entry Count | Minimum Memory |
|---|---|---|
| Small | Less than 600,000 | 512 MB |
| Medium | 600,000 to 2,000,000 | 1 GB |
| Large | Greater than 2,000,000 | 2 GB |

Going back to our example of Acme Corporation, the number of entries in the directory are close to 1,000,000 (1 million). Oracle Corporation recommends choosing the 2 GB option in order to maximize performance.

## Network Requirements

The network is rarely a bottleneck in most installations. However serious consideration must be given to it during the capacity planning stage. If the clients do not get adequate network bandwidth to send and receive messages from Oracle Internet Directory, the overall throughput will seem to be very low. For example, if we have configured Oracle Internet Directory to service 800 search operations per second, but the computer running the Oracle directory server is only accessible through a 10 Mbps network (10-Base-T switched ethernet), and we have only 60 percent of the bandwidth available, then the clients will only see a throughput of 600 search operations a second (assuming each search operation causes 1024 bytes to be transferred on the network). The following table shows the maximum possible throughput (in operations per second) for two types of operations (one requiring a transfer of 1024 bytes the other requiring a transfer of 2048 bytes) for two types of networks, 10 Mbps & 100 Mbps, at different rates of bandwidth availability:

In some cases, it may also be important to consider the network latency of sending a message from a client to the Oracle directory server. In some WAN implementations, the network latencies may become as high as 500 milliseconds, which may cause the clients to time out for certain operations. In summary, given a range of networking options, the preferred choice should always be for highest bandwidth, lowest latency network.

Going back to the example of Acme Corporation, their peak usage rate is 936,000 lookups per hour which results in an equivalent number of lookup operations to the directory. This requires about 260 directory operations per second. Assuming that each operation results in a transfer of 2 KB of data on the network, this would imply that we should have a 100 Mbps network or at least 60 percent bandwidth available on a 10 Mbps network. Since the 100 Mbps network will typically have a lower latency, we will chose that over the 10 Mbps network.

# CPU Requirements

This section contains these topics:



## CPU Configuration

The CPU sizing for Oracle Internet Directory is directly a function of the user workload. The following factors will determine CPU configuration:

- The number of concurrent operations you want to support. This will be directly dependent on the number of users performing operations simultaneously.
- The acceptable latency of each operation. For example, in an email application, a latency per operation of 100 milliseconds might be desirable, but in most cases a latency of 500 milliseconds might still be acceptable.

CPU resources can be added to a system as the workload increases, but these additions seldom bring linear scalability to all operations since a lot of operations are not purely CPU bound. We classify the processing power of a computer by a performance characteristic that is commonly available from all vendors, namely, SPECint_rate95 baseline. This number is derived from a set of integer tests and is available from all system vendors as well as the SPEC web site (`www.spec.org`).

---

**Note:** SPECint_rate95 should not be confused with the regular SPECint95 performance number. The SPECint95 performance number gives an idea of the integer processing power of a particular CPU (for systems with multiple CPUs, this number is typically normalized). The SPECint_rate95 gives the integer processing power of an entire system without any normalization.

---

Because Oracle Internet Directory makes efficient use of multiple CPUs on an SMP computer, we chose to categorize computers based on their SPECint_rate95 numbers. Even within SPECint_rate95 we chose the baseline number as opposed to the commonly advertised result. This is because the commonly advertised result is actually the peak performance of a computer, whereas the baseline number represents the performance in normal circumstances.

## Rough Estimates of CPU Requirements

Since Oracle Internet Directory is typically co-resident with the Oracle8*i* database, we recommend at least a two-CPU system. We give the following rough estimates based on the level of usage of Oracle Internet Directory:

| Usage | Num CPUs | SPECint_rate95 baseline | System |
|---|---|---|---|
| Departmental | 2 | 60 to 200 | Compaq AlphaServer 8400 5/300 (300Mhz x 2) |
| Organization wide | 4 | 200 to 350 | IBM RS/6000 J50 (200MHz x 4) |
| Enterprise wide | 4+ | 350+ | Sun Ultra 450 (296 MHz x 4) |

## Detailed Calculations of CPU Requirements

It is difficult to determine the CPU requirements for all operations at a given deployment site since the amount of CPU consumed depends upon several factors, such as:

- The type operation: base search, subtree search, modify, add etc.
- If SSL mode is enabled or not. SSL consumes an additional 15 to 20 percent of CPU resources.
- The number of entries returned for a search
- The number of access control policies that need to be checked as part of a search

In most of the cases, except SSL, we can expect that there is a large latency between the Oracle Internet Directory server process and the database. When a thread in the Oracle Internet Directory server process is waiting for the database to respond, other threads within the Oracle Internet Directory server process can be put to work by other client requests needing LDAP server specific processing. As a result, for any mix of operations, one can always come up with a combination of concurrent clients and Oracle Internet Directory server processes that will result in 100 percent CPU utilization. In this case, the CPU becomes the bottleneck.

Given this fact, we have taken the operation that consumes the smallest number of CPU cycles: a base search and estimated the number of concurrent operations at which we peaked on CPU usage on various computers. We then correlated this to SPECint_rate95 baseline number of the computers. With this correlation, given a certain amount of concurrency on the user load, one can find a lower bound on the processing power required by Oracle Internet Directory. The following formula gives the concurrency to SPECint_rate95 baseline number for this release of Oracle Internet Directory:

```
SPECint_rate95 baseline = 6.0 * (concurrent base search operations)
```

For example, if we need a computer that is capable of handling 50 concurrent base search operations before saturating the CPU, we would require a computer that has a SPECint_rate95 baseline rating of about 300.

Taking this number as the baseline, we can find the CPU requirements of other operations if we express them as some factor of the base search operations. The following factors may be used in addition to others:

- If using SSL mode, multiply CPU requirements by a factor of 1.2.
- If one is fetching a lot of entries in each search, multiply CPU requirements by a factor of (1 + 0.2* *num_entries_per_search*).
- Incorporate a factor of safety of 20 percent to 30 percent (multiply by 1.2 to 1.3).

Going back to our example of Acme Corporation, let us assume that we want adequate CPU resources to support about 100 concurrent operations. Assuming that each search returns 1.5 entries, and adding a factor of safety of 20 percent, our preliminary estimate of the CPU requirements would be:

```
SPECint_rate95 baseline=6.0*100*(1 + 0.2*1.5)*1.2 = 600*1.3*1.2 = 936
```

Looking at the available systems from the SPEC web site (`www.spec.org`) we can see that the following computer configurations would be the smallest configurations that should be considered.

The next table shows some of the computers that Acme Corporation can consider using for Oracle Internet Directory.

| Company | Model | CPUs | CPU type | SPECint95_rate baseline |
|---|---|---|---|---|
| Sun Microsystems | ES 4002 | 12 | 250MHz UltraSPARC II | 943 |
| Siemens Nixdorf | RM600 Model E60 | 8 | 250 MHz R10000 | 970 |
| Hewlett-Packard | HP SPP1600 | 32 | 120 MHz PA-RISC 7200 | 996 |
| SGI | Origin2000 | 8 | 250 MHz MIPS R10000 | 1001 |
| Data General Corporation | AViiON AV 20000 | 16 | Pentium Pro (200 MHz) | 1007 |
| Sun Microsystems | Sun Enterprise 3500 | 8 | 400MHz UltraSPARC II | 1011 |
| Sun Microsystems | Sun Enterprise 3500 | 8 | 400MHz UltraSPARC II | 1030 |
| Hewlett-Packard | HP 9000 Model N4000 | 4 | 440 MHz PA-RISC 8500 | 1093 |
| Hewlett-Packard | HP 9000 Model T600 | 12 | 180MHz PA-RISC 8000 | 1099 |
| Siemens AG | RM600 Model E80 | 8 | 285 MHz R12000 | 1103 |
| Compaq Corporation | AlphaServer 8400 5/440 | 12 | 437 MHz 21164 | 1146 |
| Compaq Corporation | AlphaServer 8400 5/625 | 8 | 612 MHz 21164 | 1153 |
| SGI | origin2000 | 16 | 195 MHz MIPS R10000 | 1182 |
| Sun Microsystems | Sun Enterprise 4000 | 12 | 336MHz UltraSPARC II | 1211 |

# Summary of Capacity Plan for Acme Corporation

In the preceding sections, we have described various components involved in capacity planning and have also shown how each of them would apply to an Oracle Internet Directory deployment at a hypothetical company named Acme Corporation. In this section we give a quick summary of all of the recommendations made. Following were the initial assumptions:

- Overall directory size: 3,200,000 entries (3.2 million)
- Number of users: 40,000
- Type of application: IMAP messaging
- Peak search rate: 260 searches / sec
- Concurrent usage rate for best CPU utilization: 100

Based on the above requirements and further assumptions, we developed the following recommendations:

- Disk space: 7 GB to 11 GB
- Memory: 2 GB
- Network: 100 Base-T
- CPU: something that has a SPECint_rate95 of at least 936.

Several simplifying assumptions were made so that the sizing calculations could be more intuitive.

# 15

# Tuning

Once capacity planning as described in Chapter 14, "Capacity Planning" is complete, and the necessary hardware acquired, it is important to perform some test runs to figure out if the hardware and software combination is yielding the desired levels of performance. This chapter gives guidelines for tuning an Oracle Internet Directory installation. This chapter contains these topics:

- About Tuning
- Tools for Performance Tuning
- CPU Usage Tuning
- Memory Tuning
- Disk Tuning
- Database Tuning
- Performance Troubleshooting

# About Tuning

The two main performance metrics for any installation of Oracle Internet Directory are:

- The average latency of individual operations at peak load

  This is the time for each operation to complete.

- The overall throughput of Oracle Internet Directory expressed in operations per second at peak load

  This is the rate at which an instance of Oracle Internet Directory is capable of completing client operations

If the performance tests yield poor results, the performance problems may be identified and fixed using the information provided in the following sections.

# Tools for Performance Tuning

Knowledge of the following tools is recommended for Solaris and most other UNIX operating systems:

| Tool | Description |
|---|---|
| top | Displays the top CPU consumers on a system |
| vmstat | Shows running statistics on various parts of the system including the Virtual Memory Manager |
| mpstat | Shows an output similar to vmstat but split across various CPUs in the system. This is available on Solaris only. |
| iostat | Shows the disk I/O statistics from various disk controllers |

Knowledge of the following tools is recommended for Windows NT:

| Tool | Description |
|---|---|
| Windows NT Performance Monitor | Provides a customized view of the events in the system |
| Windows NT Task Manager | Provides a high level output (like 'top' on UNIX) of the major things happening in the system. |

Knowledge of the following tools is recommended for Oracle8*i*:

- `utlbstat.sql` and `utlestat.sql`
- The ANALYZE function in the DBMS_STATS package

**See Also:**

- *Oracle8i Reference* for information about `utlbstat.sql` and `utlestat.sql`
- *Oracle8i Concepts* for information about the ANALYZE function in the DBMS_STATS package

In addition to the operating system tools, the LDAP applications being used in a customer environment must be able to provide latency and throughput measurement.

In addition, the Database Statistics Collection Tool (oidstats.sh), located at `$ORACLE_HOME/ldap/admin`, is provided to analyze the various database 'ods' schema objects to estimate the statistics.

**See Also:** "OID Database Statistics Collection Tool Syntax" on page A-37

## CPU Usage Tuning

The CPU is perhaps the most important resource available for any software. While Chapter 14 gives a rough estimate of the required CPU horsepower for a given application load, sometimes insufficient tuning can cause inefficient use of the CPU resources. Consider tuning CPU resources if either of the following cases is true:

- At peak loads the CPU is 100 percent utilized.
- At peak loads the CPU is underutilized, there is a significant amount of idle time in the system, and this idle time cannot be eliminated at even higher loads.

Internal benchmarks show that Oracle Internet Directory performs best when approximately 70 to 75 percent of the CPU resources are consumed by Oracle Internet Directory processes, and the remaining (about 25 to 30 percent) are consumed by the Oracle foreground processes corresponding to the database connections. While monitoring CPU usage, it is also important to monitor the percentage of time spent in the system space compared to user space. Internal benchmarks show best throughput numbers at about 85 percent user and 15 percent system time.

This section contains these topics:

- Tuning CPU for Oracle Internet Directory Processes
- Tuning CPU for Oracle Foreground Processes
- Taking Advantage of Processor Affinity on SMP Systems
- Other Alternatives for a CPU Constrained System

## Tuning CPU for Oracle Internet Directory Processes

The demands placed by Oracle Internet Directory processes on the CPU can be controlled by the ORCLSERVERPROCS and ORCLMAXCC parameters. This table lists suggested values for these parameters for various client loads:

| Parameters | 500 Concurrent LDAP Clients | 1000 Concurrent LDAP Clients | 1500 Concurrent LDAP Clients | 2000 Concurrent LDAP Clients |
|---|---|---|---|---|
| Server processes ORCLSERVERPROCS | 10 to 15 | 20 to 30 | 30 to 40 | 40 to 60 |
| Database connections ORCLMAXCC | 10 to 15 | 15 to 20 | 15 to 20 | 15 to 20 |

If we take the example of 500 concurrent clients, a value of 10 for ORCLSERVERPROCS with a value of 15 for ORCLMAXCC will result in the following configuration:

- There will be ten server processes created.
- Each server process will spawn fifteen worker threads that will do the actual work.
- Each server process will also maintain a pool of sixteen database connections (15+1) that will be shared among the worker threads.

### Tuning Oracle Internet Directory Processes When CPU Is 100 Percent Utilized

If the CPU usage of the system is at 100 percent, further tuning of the Oracle Internet Directory processes should be considered if both of the following conditions are met:

- At peak loads, Oracle Internet Directory processes consume more than 70 percent of all available CPU resources.
- At peak loads, the overall percentage of time spent in the 'system' or 'kernel' space is greater than 20 percent, and the percentage of time spent in the 'user' time is less than 80 percent.

This condition indicates that the system has too manyOracle Internet Directory server processes and database connections configured. This results in several processes or threads contending for the same CPU resources. As a result, the computer wastes a great deal of time context-switching among runnable tasks. To avoid this, one must systematically decrease the values of ORCLSERVERPROCS and ORCLMAXCC until the best performance for the peak load is achieved and the system and user time are split up as follows:

- User time: 85 percent or higher
- System time: 15 percent or lower

### Tuning Oracle Internet Directory Processes When CPU Is Under-Utilized

If the CPU usage at peak loads is not at 100 percent and the system is idle for a large percentage of the time (that is, more than 5 percent), this indicates that Oracle Internet Directory processes are under-configured and are not making the best utilization of the CPU resources. To solve this problem, one must systematically increase the values of ORCLSERVERPROCS and ORCLMAXCC until the CPU utilization reaches 100 percent and the system and user time are split up as follows:

- User time: 85 percent or higher
- System time: 15 percent or lower

## Tuning CPU for Oracle Foreground Processes

Tuning of CPU resources for Oracle Foreground processes should be considered only if both of the following conditions are met:

- The CPU usage is close to 100 percent at peak loads.
- Oracle foreground processes consume more than 30 percent of all available CPU resources.

If Oracle foreground processes are consuming excessive CPU, it implies that the queries that Oracle Internet Directory is making against the database are using too many CPU cycles. Although there is very little control available to the users on the types of underlying operations performed by the database, the following should be attempted:

- Database statistics on all of the tables and indices associated with the ODS user on the database must be collected using the ANALYZE command. This helps the cost based optimizer make better execution plans for the queries generated by Oracle Internet Directory.
- If the ANALYZE fails to produce better results, and the LDAP queries used have a lot of filters in them, then a simple reorganization of the order in which the filters are specified (with the most specific filter in the beginning and the most generic filter at the end) helps reduce the CPU consumption of the Oracle foreground processes.

## Taking Advantage of Processor Affinity on SMP Systems

Several Symmetric Multi-Processor (SMP) systems offer the capability to bind a particular process to a particular CPU. While it is generally a good idea not to bind any process to any processor, it may improve performance if the following conditions are met:

- The CPU utilization of the entire system is close to 100 percent.
- There are more than two CPUs on the computer.
- Oracle Internet Directory processes consume around 70 to 75 percent of the CPU resources.
- The database processes consume around 25 to 30 percent of the CPU resources.

Under the conditions noted above, allowing the database foreground process to run on any CPU can potentially cause many hardware cache misses for other tasks. This is because the database processes need to reference a large amount of data as part of their regular execution, and this often exceeds the limits of L2 caches available on

most systems. As a result, when the database process executes on a CPU, most of L2 cache contains pages from the **System Global Area (SGA)**. If a task switch occurs and an Oracle Internet Directory process is activated, all of its fetches from memory will be much slower because the task preceding it on the processor dirtied the L2 cache.

Restricting all of the Oracle foreground processes to execute on only one processor avoids many of the cache misses for Oracle Internet Directory processes. This, in turn, improves the overall performance.

### Other Alternatives for a CPU Constrained System

If none of the tips stated in the preceding sections solve CPU related performance problems, the following options are available:

- Upgrade the processing power of the computer, that is, add more CPUs or replace slower CPUs with faster ones.
- Keep the Oracle directory server and the associated Oracle8*i* database on separate computers.

## Memory Tuning

After the CPU, memory is the next most important thing to tune. The primary consumer of memory in an Oracle Internet Directory installation is the Oracle8*i* database. The SGA of the back-end database must be made as large as possible while leaving room for Oracle Internet Directory and Oracle processes to operate their private stacks and heaps. This section provides some details on determining various components of the SGA.

This section contains these topics:

## Tuning the System Global Area (SGA) for Oracle8*i*

The SGA should be sized based on the available physical memory on the system running Oracle8*i*.

> **See Also:** *Oracle8i Performance Guide and Reference* for more information on determining appropriate sizes for the SGA. This book tells how to ensure that the SGA size does not cause increased paging swapping activity. The latter is very detrimental to performance.

Once the available size of the SGA is determined, two primary tuning items need to be considered:

- Size of the shared pool
- Size of the buffer cache

An initial estimate for the shared pool size is .5 MB per concurrent database connection determined above.

If this estimate consumes more than 30 percent of the total SGA, use 30 percent of the total SGA instead.

Divide 60 percent of the remaining available SGA size by the block size for the database and use this value for the number of DB_BLOCK_BUFFERS. Both of these values should be initial estimates and can be refined using BSTAT/ESTAT and other RDBMS monitoring tools to determine more accurate sizes for best performance.

## Other Alternatives for a Memory-Constrained System

If there is insufficient memory to run both the database and the Oracle directory server on the same computer, then one can put the database on a different computer.

# Disk Tuning

Balancing Disk I/O is an important consideration in overall RDBMS, and hence Oracle Internet Directory performance. Typically, one can maximize the I/O throughput by using one or more of the following techniques:

- Striping logical volumes so that the I/O operations use multiple disk spindles
- Putting different tablespaces in different logical and physical disk volumes
- Distributing the disk volumes on multiple I/O controllers

> **See Also:** *Oracle8i Performance Guide and Reference* for general information about balancing and tuning disk I/O

This section contains these topics:



## Balancing Tablespaces

The Oracle Internet Directory schema is distributed among several tablespaces at installation time for ease of maintenance and performance. Each tablespace contains a grouping of Oracle Internet Directory schema objects appropriate for co-location on disk storage. As available, it is also beneficial to distribute the following objects onto separate logical disks.

> **See Also:** "RAID" on page 15-10 for more discussion about logical disks

Separate the following:

- OLTS_ATTRSTORE and OLTS_IND_ATTRSTORE

  Separating the attribute store table from its index
- OLTS_CT_DN and OLTS_IND_CT_DN

  Separating the DN catalog from its index
- OLTS_xxxx and OLTS_IND_xxxx

  (Empirically, separate the storage tablespace from the associated index)
- OLTS_IND_ATTRSTORE and OLTS_IND_CT_DN

Alternating the attribute store and DN catalog indexes. This helps even if there are only two logical disks available (one containing OLTS_CT_DN and OLTS_IND_ATTRSTORE and the other containing OLTS_IND_CT_DN and OLTS_ATTRSTORE)

### RAID

The information on balancing tablespaces is given in terms of separating Oracle Internet Directory tablespaces onto different logical drives. This assumes that a 'logical drive' is manifested on a separate disk or set of disks from other 'logical drives', and thus represents a division among disks for I/O. (Two logical drives on the same physical disk media do not really provide the same combined I/O throughput of two logical drives located on different physical media.) If a logical drive can be manifest on a striped or RAID disk subsystem, then this may increase the I/O capacity of that logical drive. However, the tablespace locations considered earlier remain applicable when considering, for instance, different logical drives of a volume manager.

## Database Tuning

This section describes the other tunable parameters available to an Oracle Internet Directory installation.

The following table gives a quick overview of the recommended values of RDBMS parameters for various client loads. These parameters are configurable in the initialization parameter file.

| Parameters | 500 Concurrent LDAP Clients | 1000 Concurrent LDAP Clients | 1500 Concurrent LDAP Clients | 2000 Concurrent LDAP Clients |
|---|---|---|---|---|
| Open cursors | 100 | 100 | 100 | 100 |
| Sessions | 225 | 600 | 800 | 1200 |
| Database block buffers | 200 to 250 MB | 200 to 250 MB | 200 to 250 MB | 200 to 250 MB |
| Database block size | 8192 | 8192 | 8192 | 8192 |
| Shared pool size | 30 to 40 MB | 30 to 40 MB | 30 to 40 MB | 30 to 40 MB |
| Processes | 400 | 800 | 1000 | 1500 |

This section describes each of the RDBMS tunable parameters in more detail. It contains these topics:



## Required Parameter

Configure the OPEN_CURSORS parameter as follows:

```
OPEN_CURSORS=100
```

The Oracle8*i* default of 50 or so is too small to accommodate Oracle Internet Directory server cursor cache. Note that this value is not dependent on other Oracle Internet Directory server parameters, such as # SERVERS and # WORKERS. The value of 100 is sufficient for any size DIT.

## Parameters Dependent on Oracle Internet Directory Server Configuration

Configure the SESSIONS parameter as follows:

```
PROCESSES = (# OID server processes per instance) x
            (# DB Connections per server + 1) x
            (# of OID instances) + 20
SESSIONS = 1.1 *PROCESSES + 5
```

Each Oracle Internet Directory server process requires a number of concurrent database connections equal to the number of worker threads configured for that server plus one. The total number of concurrent database connections allowed must therefore include this number per server, per instance. The additional 20 connections added to the parameter value accounts for the Oracle background processes plus other Oracle Internet Directory processes such as OID Monitor, OID Control, Oracle directory replication server, and bulk tools.

### Using Multi-Threaded Server (MTS)

Depending on the total number of concurrent database connections required, and as determined by the setting for the SESSIONS parameter, enabling MTS may help balance overall system load better. If the total number of concurrent database connections required is over 300, then configure MTS. One shared server should be configured for every 10 database connections required.

**Note:** The number of requires concurrent database connections depends on the hardware selected. See *Net8 Administrator's Guide* and *Oracle8i Administrator's Guide* for further information about MTS configuration.

### SGA Parameters Dependent on Hardware Resources

The main parameters that contribute to the SGA are discussed in "Memory Tuning" on page 15-7. The following are a few more parameters that may be tuned:

- Sort area

  Set to 262144 (256k) to ensure sufficient sort area available to prevent on-disk sorts.

- Redo Log Buffers

  Set to 32768 (32k) as an initial estimate. If log write performance becomes a performance problem, use a large enough value to make sure (redo log space requests / redo entries) > 1/5000 to prevent the LGWR process from falling behind. This overall has little size effect on the variable SGA size, so making this a little bit too large should not be a problem.

## Performance Troubleshooting

This section gives some quick pointers for common performance related problems.

If LDAP search performance is poor, make sure that:

- The attributes on which the search is being made are indexed
- Schema associated with the 'ODS' user is 'ANALYZED'

  For searches involving multiple filter operands, make sure that the order in which they are given goes from the 'most specific' to the 'least specific'. For example, `&(l=Chicago)(state=Illinois)(c=US)` is better than `&(c=US)(state=Illinois)(l=Chicago)`.

If LDAP add/modify performance is poor, make sure that:

- There are enough redo-log files in the database
- There are enough rollback segments in the database
- The schema associated with the 'ODS' user is 'ANALYZED'

# 16

# High Availability And Failover

This chapter discusses the high availability and failover features and deployment guidelines for Oracle Internet Directory. It contains these topics:

# About High Availability and Failover for Oracle Internet Directory

Oracle Internet Directory is designed to address the deployment needs of mission critical applications requiring a high degree of system availability. To achieve a high degree of availability, all components in the system must facilitate redundancy, and all interfaces must facilitate failure recognition and recovery, called **failover**. In addition, integration of application independent network failover capabilities in the overall deployment is also essential to achieve overall system availability.

Oracle products are commonly targeted for high availability environments and hence necessary capabilities are built into all layers of the Oracle technology stack described on page 16-2. Typically, it is not necessary to employ every failover capability in every component. This chapter describes the availability and failover features of various components in the Oracle Internet Directory / technology stack, and provides guidelines for exploiting them optimally for typical directory deployment.

# Oracle Internet Directory and Oracle8*i* Technology Stack

Figure 16–1 gives an overview of the various components of the Oracle Internet Directory stack. Stack communication between separate computers occurs by passing information from one node to the other through several layers of code. Information descends through layers on the client side. It is then packaged for transport across a network medium. The information then proceeds up the stack on the server side where it is translated and understood by the corresponding layers.

***Figure 16–1 Oracle Internet Directory/Oracle8i Technology Stack***

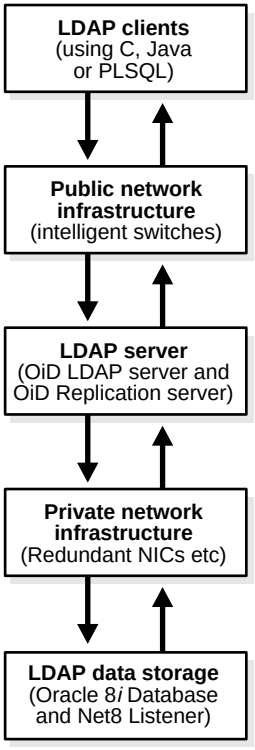


You can build sufficient fault tolerance mechanisms into each of the layers to ensure maximum availability of the product. In the following sections we describe some of the high availability options available to our customers in each of the layers shown above.

# Failover Options on Clients

Incorporating enough intelligence in the clients so that they can failover to alternate Oracle directory servers in case the primary Oracle directory server fails is a good option in some cases. This requires the clients to cache alternate server information and use it upon recognizing connectivity loss. This method of guaranteeing availability is viable only for deployments in which one has full control over the type of clients accessing the directory.

This section contains these topics:

- Alternate Server List from User Input
- Alternate Server List from the Oracle Internet Directory Server

## Alternate Server List from User Input

The clients can be designed to take input from the user on the list of alternate Oracle directory servers so that the clients can automatically failover in the event of a failure of the primary server. However, as the number of clients increases, this option would not scale very well in terms of administration of client installations.

## Alternate Server List from the Oracle Internet Directory Server

Oracle Internet Directory supports a DSE root attribute called `AlternateServers`. This is an LDAP Version 3 standard attribute and is to be maintained by the directory administrator. It is expected to have references to other Oracle directory servers in the system with the same set of naming contexts as that of the local server. When connectivity to the local server is lost, clients have the option of accessing one of the servers listed in this attribute. This option requires explicit administrative action to maintain this attribute.

**See Also:**

- "Managing Attributes by Using Oracle Directory Manager" on page 6-17 and "Managing Attributes by Using Command Line Tools" on page 6-28 to set the `AlternateServers` attribute

## Failover Options in the Public Network Infrastructure

The network used to access Oracle Internet Directory services is called the Public Network Infrastructure. Providing network level load balancing and failover measures (connection re-direction) in the Public Network Infrastructure are highly recommended since these measures provide a high degree of flexibility and transparency to the application clients.

If the Oracle Internet Directory services are accessed from the Internet, this would include a couple of high speed links (T1 to T3) and an intelligent TCP/IP level connection re-director. If the Oracle Internet Directory services are accessed from an Intranet, this would include high speed LAN connections to the server computers running the Oracle directory server and an intelligent TCP/IP level connection re-director. In both cases, there would be more than one computer serving LDAP requests so that failure of one Oracle directory server computer would not affect availability.

Figure 16–2 illustrates a typical Internet deployment of Oracle Internet Directory with network-level failover enabled.

***Figure 16–2 Network-Level Failover***

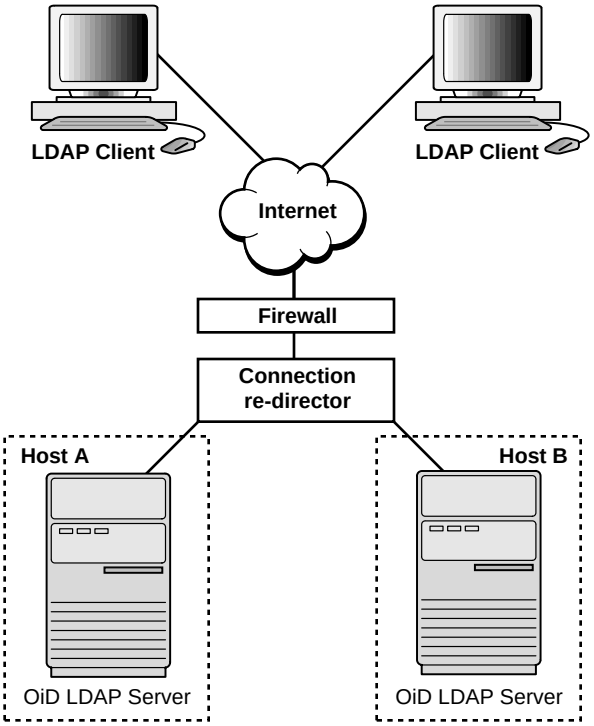


In Figure 16–2, the Oracle directory servers (OiD LDAP Servers) can be connected to either the same back-end database or different back-end databases. In this deployment, network-level connection redirection can be accomplished by both hardware and software solutions.

This section contains these topics:

- Hardware-Based Connection Redirection
- Software-Based Connection Redirection

### Hardware-Based Connection Redirection

Hardware-based connection redirection technology is available from several vendors. These redirection devices connect directly to the Internet and can route requests among several server computers. They can also detect computer failures and stop routing requests to the failed computer. This feature guarantees that new connections from clients will not be routed to a failed computer. When a computer comes back, the device detects it and starts routing new requests to it. These devices also perform some load balancing, which makes sure that client requests are uniformly distributed.

Some of the vendors providing hardware based re-direction technologies are:

- Accelar Server Switches from Nortel Networks
- Local Director from Cisco
- BIG/ip from F5 Labs Inc.
- Hydra from HydraWEB Technologies
- Equalizer from Coyote Point Systems

### Software-Based Connection Redirection

The software-based solutions essentially work in the same manner as their hardware counterparts. Some of the currently available solutions include Dispatch from Resonate and Network Dispatcher from IBM.

## Availability and Failover Capabilities in Oracle Internet Directory

Multi-master replication makes it possible for the directory system to be available for both access and updates at all times, as long as at least one of the nodes in the system is available. When a node comes back online after a period of unavailability, replication from the existing nodes will resume automatically and cause its contents to be synchronized transparently.

Any directory system with high availability requirements should always employ a network of replicated nodes in multi-master configuration. A replica node is recommended for each region that is separated from others by a relatively low speed or low bandwidth network segment. Such a configuration, while allowing speedy directory access to the clients in the same region, also serves as a failover arrangement during regional failures elsewhere.

# Failover Options in the Private Network Infrastructure

The Private Network Infrastructure is the network used by Oracle Internet Directory and its back-end components to communicate with each other. In cases where Oracle Internet Directory is deployed on the Internet, Oracle Corporation recommends that this network be physically different from the network used to serve client requests. In cases where Oracle Internet Directory is deployed over an Intranet, the same LAN may be used, but Oracle Internet Directory components should have dedicated bandwidth with the help of a network switch. Because Oracle Internet Directory depends on the Private Network Infrastructure for its communications, you must take adequate precautions to guarantee availability in the event of failures in the Private Network. Some of the options available in this area are:

- IP Address Takeover (IPAT)
- Redundant Links

## IP Address Takeover (IPAT)

IP address takeover feature is available on many commercial clusters. This feature protects an installation against failures of the Network Interface Cards (NICs). In order to make this mechanism work, installations must have two NICs for each IP address assigned to a server. Both the NICs must be connected to the same physical network. One NIC is always active while the other is in a standby mode. The moment the system detects a problem with the main adapter, it immediately fails over to the standby NIC. Ongoing TCP/IP connections are not disturbed and as a result clients do not notice any downtime on the server.

## Redundant Links

Since all networks (with the exception of wireless networks) are comprised of wires going from one location to the other, there is a distinct possibility that someone might unintentionally disconnect a wire that is used to link a client computer to a server computer. If you want to take such precautions, use NICs and hubs/switches that come with the capability to use redundant links in case of a link level failure.

# High Availability Deployment Examples

In Figure 16–3, the database and Oracle directory server (OiD LDAP Server) are co-resident on the same computer. Changes made on one Oracle directory server instance are reflected on the second Oracle directory server instance through multimaster replication. When a failure of the Oracle directory server or database server on a particular node occurs, it is elevated to a computer failure so that the connection redirector will stop handing off connections to the computer on which there was a failure.

***Figure 16–3 Deployment Example (Two Oracle Internet Directory Nodes in Replication)***

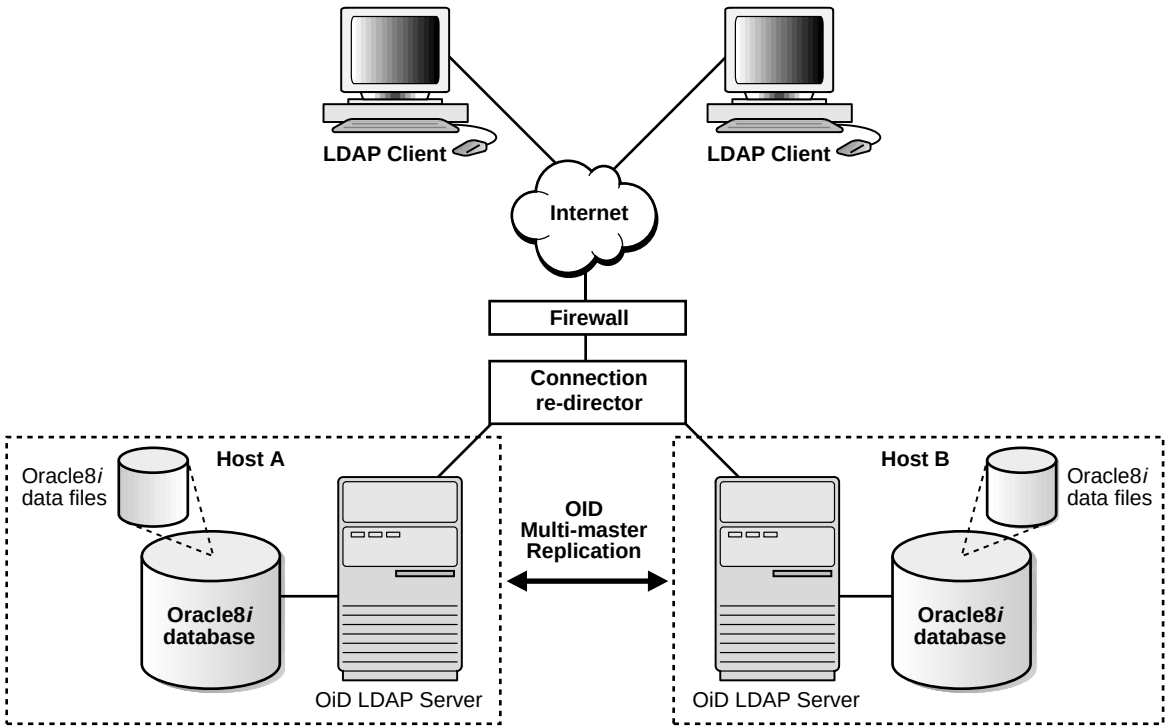

As Figure 16–4 illustrates, each of the regions can be set up with two Oracle Internet Directory nodes replicating between each other. This configuration is typical of global directory networks deployed by large enterprises where each of the regions above could potentially represent a continent or a country.

***Figure 16–4 Deployment Example 2***

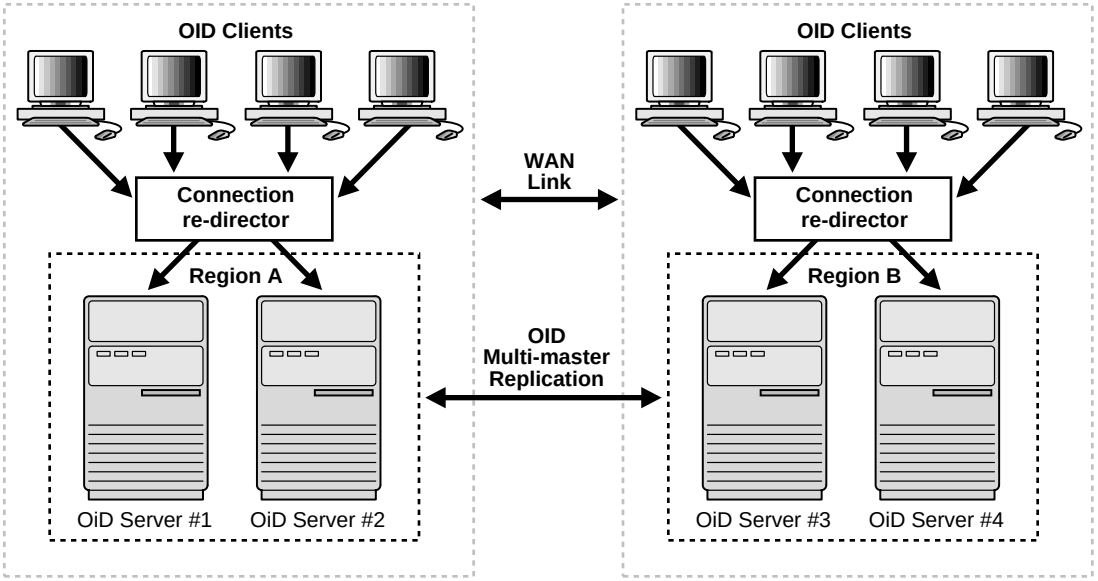

# Part IV

## Appendixes

Part IV contains the following appendixes:

- Appendix A, "Syntax for LDIF and Command Line Tools"
- Appendix B, "Adding a DSA Using the Database Copy Procedure"
- Appendix C, "Using Oracle Wallet Manager"
- Appendix D, "Using Access Control Directive Format"
- Appendix E, "Schema Elements"
- Appendix F, "Migrating Data from Other LDAP-Compliant Directories"
- Appendix G, "Troubleshooting"

# A

# Syntax for LDIF and Command Line Tools

This appendix provides syntax, usage notes, and examples for **LDAP Data Interchange Format (LDIF)** and LDAP command line tools. It contains these topics:

# LDAP Data Interchange Format (LDIF) Syntax

The standardized file format for directory entries is as follows:

```
dn: distinguished_name
attribute_type: attribute_value
.
.
.
objectClass: object_class_value
.
.
.
```

| Property | Value | Description |
|---|---|---|
| dn: | *RDN,RDN,RDN,* ... | Separate RDNs with commas. |
| *attribute*: | *attribute_value* | This line repeats for every attribute in the entry, and for every attribute value in multi-valued attributes. |
| objectClass: | *object_class_ value* | This line repeats for every object class. |

The following example shows a file entry for an employee. The first line contains the DN. The lines that follow the DN begin with the mnemonic for an attribute, followed by the value to be associated with that attribute. Note that each entry ends with lines defining the object classes for the entry.

```
dn: cn=Suzie Smith,ou=Server Technology,o=Acme, c=US
cn: Suzie Smith
cn: SuzieS
sn: Smith
email: ssmith@us.Acme.com
telephoneNumber: 69332
photo:/ORACLE_HOME/empdir/photog/ssmith.jpg
objectClass: organizational person
objectClass: person
objectClass: top
```

The next example shows a file entry for an organization.

```
dn: o=Acme,c=US
o: Acme
ou: Financial Applications
objectClass: organization
objectClass: top
```

## LDIF Formatting Notes

A list of formatting rules follows. This list is not exhaustive.

- All mandatory attributes belonging to an entry being added must be included with non-null values in the LDIF file.

  **Tip:** To see the mandatory and optional attribute types for an object class, use Oracle Directory Manager. See "Viewing Properties of Object Classes by Using Oracle Directory Manager" on page 6-9.

- Non-printing characters and tabs are represented in attribute values by base-64 encoding.
- The entries in your file must be separated from each other by a blank line.
- A file must contain at least one entry.
- Lines can be continued to the next line by beginning the continuation line with a space or a tab.
- Add a blank line between separate entries.
- Reference binary files, such as photographs, with the absolute address of the file, preceded by a forward slash ("/").
- The DN contains the full, unique directory address for the object.
- The lines listed after the DN contain both the attributes and their values. DNs and attributes used in the input file must match the existing structure of the DIT. Do not use attributes in the input file that you have not implemented in your DIT.
- Sequence the entries in an LDIF file so that the DIT is created from the top down. If an entry relies on an earlier entry for its DN, make sure that the earlier entry is added before its child entry.
- When you define schema within an LDIF file, insert a white space between the opening parenthesis and the beginning of the text, and between the end of the text and the ending parenthesis.

**See Also:**

- The various resources listed in "Related Documentation" on page xxvii for a complete list of LDIF formatting rules
- "Using NLS with LDIF Files" on page 12-3

# Command Line Tools Syntax

This section tells you how to use the following tools:

- ldapadd Syntax
- ldapaddmt Syntax
- ldapbind Syntax
- ldapcompare Syntax
- ldapdelete Syntax
- ldapmoddn Syntax
- ldapmodify Syntax
- ldapmodifymt Syntax
- ldapsearch Syntax

## ldapadd Syntax

The ldapadd command line tool enables you to add entries, their object classes, attributes, and values to the directory. To add attributes to an existing entry, use the ldapmodify command, explained in "ldapmodify Syntax" on page A-13.

**See Also:** "Adding Configuration Set Entries by Using ldapadd" on page 5-11 for an explanation of using ldapadd to configure a server with an input file

ldapadd uses this syntax:

```
ldapadd [arguments] -f filename
```

where *filename* is the name of an LDIF file written with the specifications explained in the section "LDAP Data Interchange Format (LDIF) Syntax" on page A-2.

The following example adds the entry specified in the LDIF file `my_ldif_file.ldi`:

```
ldapadd -p 389 -h myhost -f my_ldif_file.ldi
```

| Optional Arguments | Descriptions |
|---|---|
| -b | Specifies that you have included binary file names in the file, which are preceded by a forward slash character. The tool retrieves the actual values from the file referenced. |
| -c | Tells ldapadd to proceed in spite of errors. The errors will be reported. (If you do not use this option, ldapadd stops when it encounters an error.) |
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry specified in *binddn*. Use this with the `-w` *password* option. |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -f *filename* | Specifies the input name of the LDIF format import data file. For a detailed explanation of how to format an LDIF file, see "LDAP Data Interchange Format (LDIF) Syntax" on page A-2. |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -K | Same as `-k`, but performs only the first step of the Kerberos bind |
| -k | Authenticates using Kerberos authentication instead of simple authentication. To enable this option, you must compile with KERBEROS defined.<br>You must already have a valid ticket granting ticket. |
| -n | Shows what would occur without actually performing the operation |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password required for one-way or two-way SSL connections |
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |

| Optional Arguments | Descriptions |
|---|---|
| -v | Specifies verbose mode |
| -w *password* | Provides the password required to connect |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |

## ldapaddmt Syntax

ldapaddmt is like ldapadd: it enables you to add entries, their object classes, attributes, and values to the directory. It is unlike ldapadd in that it supports multiple threads for adding entries concurrently.

While it is processing LDIF entries, ldapaddmt logs errors in the `add.log` file in the current directory.

ldapaddmt uses this syntax:

```
ldapaddmt -T number_of_threads -h host -p port -f filename
```

where *`filename`* is the name of an LDIF file written with the specifications explained in the section "LDAP Data Interchange Format (LDIF) Syntax" on page A-2.

The following example uses five concurrent threads to process the entries in the file `myentries.ldif`.

```
ldapaddmt -T 5 -h node1 -p 3000 -f myentries.ldif
```

**Note:** Increasing the number of concurrent threads improves the rate at which LDIF entries are created, but consumes more system resources.

| Optional Arguments | Descriptions |
|---|---|
| -b | Specifies that you have included binary file names in the data file, which are preceded by a forward slash character. The tool retrieves the actual values from the file referenced. |
| -c | Tells the tool to proceed in spite of errors. The errors will be reported. (If you do not use this option, the tool stops when it encounters an error.) |
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry is specified in *binddn*. Use this with the `-w password` option. |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)" |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -K | Same as -k, but performs only the first step of the kerberos bind |
| -k | Authenticates using Kerberos authentication instead of simple authentication. To enable this option, you must compile with KERBEROS defined.<br>You must already have a valid ticket granting ticket. |
| -n | Shows what would occur without actually performing the operation. |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password required for one-way or two-way SSL connections |
| -T | Sets the number of threads for concurrently processing entries |
| -U *SSLAuth* | Specifies SSL Authentication Mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -v | Specifies verbose mode |
| -w *password* | Provides the password required to connect |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |

## ldapbind Syntax

The ldapbind command line tool enables you to see whether you can authenticate a client to a server.

ldapbind uses this syntax:

`ldapbind [`*`arguments`*`]`

| Optional Arguments | Descriptions |
|---|---|
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry specified in *binddn*. Use this with the `-w` *password* option. |
| -E "*.character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -n | Shows what would occur without actually performing the operation |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies the wallet password required for one-way or two-way SSL connections |
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -w *password* | Provides the password required to connect |
| -W *wallet_location* | Specifies wallet location (required for one-way or two-way SSL connections) |

## ldapcompare Syntax

The ldapcompare command line tool enables you to match attribute values you specify in the command line with the attribute values in the directory entry.

ldapcompare uses this syntax:

```
ldapcompare [arguments]
```

The following example tells you whether `Person Nine`'s title is `associate`.

```
ldapcompare -p 389 -h myhost -b "cn=Person Nine, ou=EuroSInet Suite, o=IMC,
c=US" -a title -v associate
```

| Mandatory Arguments | Descriptions |
|---|---|
| -a *attribute name* | Specifies the attribute on which to perform the compare |
| -b *basedn* | Specifies the distinguished name of the entry on which to perform the compare |
| -v *attribute value* | Specifies the attribute value to compare |

| Optional Arguments | Descriptions |
|---|---|
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry is specified in *binddn*. Use this with the `-w` *password* option. |
| -d *debug-level* | Sets the debugging level. See "Setting Debug Logging Levels by Using the OID Control Utility" on page 5-23. |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -f *filename* | Specifies the input filename |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password (required for one-way or two-way SSL connections) |

| Optional Arguments | Descriptions |
|---|---|
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -w *password* | Provides the password required to connect |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |

## ldapdelete Syntax

The ldapdelete command line tool enables you to remove entire entries from the directory that you specify in the command line.

ldapdelete uses this syntax:

```
ldapdelete [arguments] "entry_DN"
```

The following example uses port 389 on a host named myhost.

```
ldapdelete -p 389 -h myhost ou=EuroSInet Suite, o=IMC, c=US"
```

| Optional Arguments | Descriptions |
|---|---|
| -D *binddn* | When authenticating to the directory, uses a full DN for the *binddn* parameter; typically used with the `-w` *password* option. |
| -d *debug-level* | Sets the debugging level. See "Setting Debug Logging Levels by Using the OID Control Utility" on page 5-23. |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -f *filename* | Specifies the input filename |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -k | Authenticates using authentication instead of simple authentication. To enable this option, you must compile with Kerberos defined.<br>You must already have a valid ticket granting ticket. |
| -n | Shows what would be done, but doesn't actually delete |

| Optional Arguments | Descriptions |
|---|---|
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password required for one-way or two-way SSL connections |
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -v | Specifies verbose mode |
| -w *password* | Provides the password required to connect. |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |

## ldapmoddn Syntax

The ldapmoddn command line tool enables you to modify the DN or RDN of an entry.

ldapmoddn uses this syntax:

```
ldapmoddn [arguments]
```

The following example uses ldapmoddn to modify the RDN component of a DN from `"cn=dcpl"` to `"cn=thanh mai"`. It uses port 389, and a host named myhost.

```
ldapmoddn -p 389 -h myhost -b "cn=dcpl,dc=Americas,dc=imc,dc=com" -R "cn=thanh
mai"
```

| Mandatory Argument | Description |
|---|---|
| -b *basedn* | Specifies DN of the entry to be moved |

| Optional Arguments | Descriptions |
|---|---|
| -D *binddn* | When authenticating to the directory, do so as the entry is specified in *binddn*. Use this with the `-w` *password* option. |

| Optional Arguments | Descriptions |
|---|---|
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -f *filename* | Specifies the input filename |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -N *newparent* | Specifies new parent of the RDN |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password required for one-way or two-way SSL connections |
| -r | Specifies that the old RDN is not retained as a value in the modified entry. If this argument is not included, the old RDN is retained as an attribute in the modified entry. |
| -R *newrdn* | Specifies new RDN |
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -w *password* | Provides the password required to connect. |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |

## ldapmodify Syntax

The ldapmodify tool enables you to act on attributes.

ldapmodify uses this syntax:

```
ldapmodify [arguments] -f filename
```

where *filename* is the name of an LDIF file written with the specifications explained the section "LDAP Data Interchange Format (LDIF) Syntax" on page A-2.

The list of arguments in the following table is not exhaustive.

| Optional Arguments | Description |
|---|---|
| -a | Denotes that entries are to be added, and that the input file is in LDIF format. |
| -b | Specifies that you have included binary file names in the data file, which are preceded by a forward slash character. |
| -c | Tells ldapmodify to proceed in spite of errors. The errors will be reported. (If you do not use this option, ldapmodify stops when it encounters an error.) |
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry is specified in *binddn*. Use this with the `-w` *password* option. |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -n | Shows what would occur without actually performing the operation. |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password required for one-way or two-way SSL connections |
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -v | Specifies verbose mode |

| Optional Arguments | Description |
|---|---|
| -w *password* | Overrides the default, unauthenticated, null bind. To force authentication, use this option with the `-D` option. |
| -W *wallet_location* | Specifies wallet location (required for one-way or two-way SSL connections) |

To run `modify`, `delete`, and `modifyrdn` operations using the `-f` flag, use LDIF for the input file format (see "LDAP Data Interchange Format (LDIF) Syntax" on page A-2) with the specifications noted below:

Always separate entries with a blank line.

Unnecessary space characters in the LDIF input file, such as a space at the end of an attribute value, will cause the LDAP operations to fail.

**Line 1:** Every change record has, as its first line, the literal `dn:` followed by the DN value for the entry, for example:

```
dn:cn=Barbara Fritchy,ou=Sales,o=Oracle,c=US
```

**Line 2:** Every change record has, as its second line, the literal "`changetype:`" followed by the type of change (`add, delete, modify, modrdn`), for example:

```
changetype:modify
```

or

```
changetype:modrdn
```

Format the remainder of each record according to the following requirements for each type of change:

- `changetype:add`

  Uses LDIF format (see "LDAP Data Interchange Format (LDIF) Syntax" on page A-2).

- `changetype:modify`

  The lines that follow this changetype consist of changes to attributes belonging to the entry that you identified in Line 1 above. You can specify three different types of attribute modifications—add, delete, and replace—which are explained next:

– **Add attribute values**. This option to changetype modify adds more values to an existing multi-valued attribute. If the attribute does not exist, it adds the new attribute with the specified values:

```
add: attribute name
attribute name: value1
attribute name: value2...
```

For example:

```
dn:cn=Barbara Fritchy,ou=Sales,o=Oracle,c=US
changetype:modify
add: work-phone
work-phone:510/506-7000
work-phone:510/506-7001
```

– **Delete values**. If you supply only the "delete" line, all the values for the specified attribute are deleted. Otherwise, if you specify an attribute line, you can delete specific values from the attribute:

```
delete: attribute name
[attribute name: value1]
```

For example:

```
dn:cn=Barbara Fritchy,ou=Sales,o=Oracle,c=US
changetype:delete
delete: home-fax
```

– **Replace values.** Use this option to replace all the values belonging to an attribute with the new, specified set:

```
replace:attribute name
[attribute name:value1 ...]
```

If you do not provide any attributes with "replace," the directory adds an empty set. It then interprets the empty set as a delete request, and complies by deleting the attribute from the entry. This is useful if you want to delete attributes that may or may not exist.

For example:

```
dn:cn=Barbara Fritchy,ou=Sales,o=Oracle,c=US
changetype:modify
replace: work-phone
work-phone:510/506-7002
```

* `changetype:delete`

  This change type deletes entries. It requires no further input, since you identified the entry in Line 1 and specified a changetype of delete in Line 2.

  For example:

  ```
  dn:cn=Barbara Fritchy,ou=Sales,o=Oracle,c=US
  changetype:delete
  ```

* `changetype:modrdn`

  The line following the change type provides the new relative distinguished name using this format:

  `newrdn: RDN`

  For example:

  ```
  dn:cn=Barbara Fritchy,ou=Sales,o=Oracle,c=US
  changetype:modrdn
  newrdn: cn=Barbara Fritchy-Blomberg
  ```

## ldapmodifymt Syntax

The ldapmodifymt command line tool enables you to modify several entries concurrently.

ldapmodifymt uses this syntax:

`ldapmodifymt -T number_of_threads [arguments] -f filename`

where *filename* is the name of an LDIF file written with the specifications explained the section "LDAP Data Interchange Format (LDIF) Syntax" on page A-2.

> **See Also:** "ldapmodify Syntax" on page A-13 for additional formatting specifications used by ldapmodifymt

The following example uses five concurrent threads to modify the entries in the file `myentries.ldif`.

```
ldapmodifymt -T 5 -h node1 -p 3000 -f myentries.ldif
```

| Optional Arguments | Descriptions |
|---|---|
| -a | Denotes that entries are to be added, and that the input file is in LDIF format. (If you are running ldapadd, this flag is not required.) |
| -b | Specifies that you have included binary file names in the data file, which are preceded by a forward slash character. |
| -c | Tells ldapmodify to proceed in spite of errors. The errors will be reported. (If you do not use this option, ldapmodify stops when it encounters an error.) |
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry is specified in binddn. Use this with the `-w` *password* option. |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -n | Shows what would occur without actually performing the operation. |
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password required for one-way or two-way SSL connections |
| -T | Sets the number of threads for concurrently processing entries |
| -U *SSLAuth* | Specifies SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -v | Specifies verbose mode |
| -w *password* | Overrides the default, unauthenticated, null bind. To force authentication, use this option with the -D option. |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |

## ldapsearch Syntax

The ldapsearch command line tool enables you to search for and retrieve specific entries in the directory.

ldapsearch uses this syntax:

```
ldapsearch [arguments] filter [attributes]
```

The *filter* format must be compliant with RFC-2254. For further information about this standard, search for the standard at: http://www.ietf.org/rfc/rfc2254.txt

Separate attributes with a space. If you do not list any attributes, all attributes are retrieved.

| Mandatory Arguments | Descriptions |
|---|---|
| -b *basedn* | Specifies base dn for search |
| -s *scope* | Specifies search scope: base, one, or sub. |

| Optional Arguments | Descriptions |
|---|---|
| -A | Retrieves attribute names only (no values) |
| -a *deref* | Specifies alias dereferencing: never, always, search, or find |
| -B | Allows printing of non-ASCII values |
| -D *binddn* | When authenticating to the directory, specifies doing so as the entry specified in *binddn*. Use this with the `-w` *password* option. |
| -d *debug level* | Sets debugging level to the level specified (see Table 5–1 on page 5-24) |
| -E "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| -f *file* | Performs sequence of searches listed in *file* |
| -F sep | Prints '`sep`' instead of '`=`' between attribute names and values |
| -h *ldaphost* | Connects to *ldaphost*, rather than to the default host, that is, your local computer. *ldaphost* can be a computer name or an IP address. |
| -L | Prints entries in LDIF format (`-B` is implied) |
| -l *timelimit* | Specifies maximum time (in seconds) to wait for ldapsearch command to complete |
| -n | Shows what would be done without actually searching |

| Optional Arguments | Descriptions |
|---|---|
| -p *ldapport* | Connects to the directory on TCP port *ldapport*. If you do not specify this option, the tool connects to the default port (389). |
| -P *wallet_password* | Specifies wallet password (required for one-way or two-way SSL connections) |
| -S *attr* | Sorts the results by attribute *attr* |
| -t | Writes to files in `/tmp` |
| -u | Includes user friendly entry names in the output |
| -U *SSLAuth* | Specifies the SSL authentication mode:<br>■ 1 for no authentication required<br>■ 2 for one way authentication required<br>■ 3 for two way authentication required |
| -v | Specifies verbose mode |
| -w *passwd* | Specifies bind passwd for simple authentication |
| -W *wallet_location* | Specifies wallet location required for one-way or two-way SSL connections |
| -z *sizelimit* | Specifies maximum number of entries to retrieve |

## Examples of ldapsearch Filters

Study the following examples to see how to build your own search commands.

**Example 1: Base Object Search** The following example performs a base-level search on the directory from the root.

```
ldapsearch -p 389 -h myhost -b "" -s base -v "objectclass=*"
```

- `-b` specifies base dn for search, root in this case.
- `-s` specifies whether the search is a base search (`base`), one level search (`one`) or subtree search (`sub`).
- `"objectclass=*"` specifies the filter for search.

**Example 2: One-Level Search** The following example performs a one level search starting at "ou=HR, ou=Americas, o=IMC, c=US".

```
ldapsearch -p 389 -h myhost -b "ou=HR, ou=Americas, o=IMC, c=US" -s one -v
"objectclass=*"
```

**Example 3: Sub-Tree Search** The following example performs a sub-tree search and returns all entries having a DN starting with "cn=Person".

```
ldapsearch -p 389 -h myhost -b "c=US" -s sub -v "cn=Person*"
```

**Example 4: Search Using Size Limit** The following example actually retrieves only two entries, even if there are more than two matches.

```
ldapsearch -h myhost -p 389 -z 2 -b "ou=Benefits,ou=HR,ou=Americas,o=IMC,c=US"
-s one "objectclass=*"
```

**Example 5: Search with Required Attributes and Attribute Options** The following example returns only the DN attribute values of the matching entries:

```
ldapsearch -p 389 -h myhost -b "c=US" -s sub -v "objectclass=*" dn
```

The following example retrieves only the distinguished name (dn) along with the surname (sn) and description (description) attribute values:

```
ldapsearch -p 389 -h myhost -b "c=US" -s sub -v "cn=Person*" dn sn description
```

The following example retrieves entries with common name (cn) attributes that have an option specifying a language code attribute option. This particular example retrieves entries in which the common names are in French and begin with the letter R.

```
ldapsearch -p 389 -h myhost -b "c=US" -s sub "cn;lang-fr=R*"
```

Suppose that, in the entry for John, no value is set for the cn;lang-it language code attribute option. In this case, the following example fails:

```
ldapsearch -p 389 -h myhost -b "c=us" -s sub "cn;lang-it=Giovanni
```

**Example 6: Searching for All User Attributes and Specified Operational Attributes** The following example retrieves all user attributes and the createtimestamp and orclguid operational attributes:

```
ldapsearch -p 389 -h myhost -b "ou=Benefits,ou=HR,ou=Americas,o=IMC,c=US" -s sub
"cn=Person*" * createtimestamp orclguid
```

The following example retrieves entries modified by Anne Smith:

```
ldapsearch -h sun1 -b "" "(&(objectclass=*)(modifiersname=cn=Anne
Smith))"
```

The following example retrieves entries modified between 01 April 2000 and 06 April 2000:

```
ldapsearch -h sun1 -b "" "(&(objectclass=*)(modifytimestamp>=20000401000000)
(modifytimestamp<= 20000406235959))"
```

> **Note:** Because `modifiersname` and `modifytimestamp` are not indexed attributes, use catalog.sh to index these two attributes. Then, restart the Oracle directory server before issuing the two previous ldapsearch commands.

**Other Examples:** Each of the following examples searches on port 389 of host sun1, and searches the whole subtree starting from the DN `"ou=hr,o=acme,c=us"`.

The following example searches for all entries with any value for the objectclass attribute.

```
ldapsearch -p 389 -h sun1 -b "ou=hr, o=acme, c=us" -s subtree "objectclass=*"
```

The following example searches for all entries that have `orcle` at the beginning of the value for the `objectclass` attribute.

```
ldapsearch -p 389 -h sun1 -b "ou=hr, o=acme, c=us" -s subtree
"objectclass=orcle*"
```

The following example searches for entries where the `objectclass` attribute begins with `orcle` and `cn` begins with foo.

```
ldapsearch -p 389 -h sun1 -b "ou=hr, o=acme, c=us" -s subtree
"(&(objectclass=orcle*)(cn=foo*))"
```

The following example searches for entries in which the common name (`cn`) is not `foo`.

```
ldapsearch -p 389 -h sun1 -b "ou=hr, o=acme, c=us" -s subtree "(!(cn=foo))"
```

The following example searches for entries in which `cn` begins with `foo` or `sn` begins with `bar`.

```
ldapsearch -p 389 -h sun1 -b "ou=hr, o=acme, c=us" -s subtree
"(|(cn=foo*)(sn=bar*))"
```

The following example searches for entries in which `employeenumber` is less than or equal to 10000.

```
ldapsearch -p 389 -h sun1 -b "ou=hr, o=acme, c=us" -s subtree
"employeenumber<=10000"
```

# Bulk Tools Syntax

This section contains these topics:

- bulkdelete Syntax
- bulkload Syntax
- bulkmodify Syntax
- ldifwrite Syntax

## bulkdelete Syntax

The bulkdelete command line tool enables you to delete a subtree efficiently. It can be used when both an Oracle directory server and Oracle directory replication servers are in operation. It uses a SQL interface to benefit performance. For this release, the bulkdelete tool runs on only one node at a time.

This tool does not support filter-based deletion. That is, it deletes an entire subtree below the root of the subtree. If the base DN is a user-added DN, rather than a DN created as part of the installation of the directory, it is included in the delete. You must restrict LDAP activity against the subtree during deletion.

The bulkdelete tool uses this syntax:

```
bulkdelete.sh -connect net_service_name -base "base_dn" -size number_of_entries
-encode "character_set"
```

| Mandatory Arguments | Descriptions |
|---|---|
| - connect *net_service_name* | Specifies the net service name to connect to the directory database<br>**See Also:** *Net8 Administrator's Guide* |
| - base "*base_dn*" | Specifies the base DN of the subtree to be deleted |

| Optional Arguments | Descriptions |
|---|---|
| -size *number_of_entries* | Specifies the number of entries to be committed as a part of one transaction. |
| -encode "*character_set*" | Native character set encoding |

## bulkload Syntax

The bulkload command line tool uses Oracle SQL*Loader to create directory entries from data residing in or created by other applications. When using bulkload, you specify any options and the input filename. The bulkload tool expects the input file to be in LDIF.

**See Also:** "LDAP Data Interchange Format (LDIF) Syntax" on page A-2.

The bulkload tool uses this syntax:

```
bulkload.sh -connect net_service_name [-check] [-generate] [-load]
   [-restore] absolute_path_to_ldif.file
```

| Mandatory Argument | Description |
|---|---|
| connect *net_service_name* | Specifies the net service name defined in the `tnsnames.ora` file.<br>**See Also:** *Net8 Administrator's Guide* |

| Optional Arguments | Descriptions |
|---|---|
| check | Checks LDAP schema for inconsistencies and for existence of duplicate DNs in the file |
| -encode "*character_set*" | Specifies native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |
| generate | Creates files suitable for loading into Oracle Internet Directory |
| load | Loads files resulting from generate phase into specified database |
| restore | Takes the operational attributes, such as `orclguid`, `creatorsname`, and `createtimestamp`, from the LDIF file rather than generating new ones. Use this argument only when the LDIF file contains operational attributes. Use this in conjunction with the `generate` and `check` arguments. |

Bulk loading must be performed when Oracle Internet Directory instances are not running.

> **See Also:** Chapter 5, "Managing an Oracle Directory Server" for instructions on stopping directory server instances

The LDIF data file path must be fully specified for check or generate operations.

### Bulk Loading Multiple Nodes in a Replicated Environment

After generating a file with the `generate` option, you can use the `load` option to load multiple computers with the identical SQL*Loader file. Do this only when creating a new replica node.

> **See Also:** "Task 6: Start the Replication Servers on All the Nodes" on page 10-18

The current version of bulkload does not allow you to specify the connection information for all of the nodes in one command.

When you load the same data into multiple nodes in a replicated network, ensure that the `orclGUID` parameter (global IDs) is consistent across all the nodes. You can accomplish this by generating the bulkload data file once only (using the `-generate` option), and then using the same data file to load the other nodes (using the `-load` option).

## bulkmodify Syntax

The bulkmodify command line tool enables you to modify a large number of existing entries in an efficient way. The bulkmodify tool supports the following:

- Subtree based modification
- A single attribute filter. For example, the filter could be `objectclass=*`, `objectclass=oneclass`, or `telephonenumber=*`.
- Attribute value addition and replacement. It modifies all matched entries in bulk.

The bulkmodify tool performs schema checking on the specified attribute name and value pair during initialization. All entries that meet the following criteria are modified:

- They are under the specified subtree.
- They meet the single filter condition.
- They contain the attribute to be modified as either mandatory or optional.

The Oracle directory server and Oracle directory replication server may be running concurrently while bulk modification is in progress, but the bulk modification does not affect the replication server. You must perform bulk modification against all replicas.

> **Note:** LDIF file based modification is not supported by bulkmodify. This type of modification requires per entry based schema checking, and therefore the performance gain over the existing ldapmodify tool is insignificant.

You must restrict user access to the subtree during bulk modification. If necessary, **ACI** restriction can be applied to the subtree being updated by bulkmodify.

You cannot use bulkmodify to add a value to single-valued attributes that already contain one value. If a second value is added, you must alter the directory schema to make that attribute multi-valued.

The bulkmodify tool uses this syntax:

```
bulkmodify -c net_service_name -b base_dn {-a|-r} attr_name -v att_value [-f
filter] [-s size]
```

| Mandatory Arguments | Descriptions |
|---|---|
| -c *net_service_name* | Specifies the net service name of the directory database<br>**See Also:** *Net8 Administrator's Guide* |
| -b *base_dn* | Specifies the base DN of the subtree to be modified |
| -a *attr_name* | Specifies the attribute name for addition |
| -r *attr_name* | Specifies the attribute name for replacement |
| -v *att_value* | Specifies the attribute value for either addition or replacement |

| Optional Arguments | Descriptions |
|---|---|
| -f *filter* | Specifies the filter to be used |
| -s *number_of_entries* | Specifies the number of entries to be committed as a part of one transaction. If not specified, default is 100. |
| -E "*character_set*" | Native character set encoding. See Chapter 12, "Managing National Language Support (NLS)". |

The filter specified with the `-f` option must contain a single attribute.

If a filter is not specified, the default filter `objectclass=*` is assumed.

There can be only one attribute name specified in the -a or `-r` option in each execution.

There can be only one value specified in the `-v` option in each execution. For example, the following bulkmodify command adds the telephone number 408-123-4567 to the entries of all employees who have Anne Smith as their manager:

```
-c my_database -b "c=US" -a telephoneNumber -v "408-123-4567 -f "manager=Anne
Smith"
```

To assure that the modified entries are read, after completing the bulkmodify procedure, restart the Oracle Internet Directory server.

## ldifwrite Syntax

The ldifwrite command line tool enables you to convert all or part of the information residing in an Oracle Internet Directory to LDIF. This makes that information available for loading into a new node in a replicated directory or into another node for backup storage. The ldifwrite tool performs a subtree search, including all entries below the specified DN, including the DN itself.

The ldifwrite tool uses this syntax:

```
ldifwrite -c net_service_name -b base_DN -f filename
```

| Mandatory Arguments | Descriptions |
|---|---|
| -c *net_service_name* | Specifies the net service name of the directory that is the source of the data, as defined in the `tnsnames.ora` file.<br>**See Also:** *Net8 Administrator's Guide* |
| -b *base_DN* | Specifies the base of the subtree to be written out in LDIF format |
| -f *filename* | Specifies the name of the LDIF file to be created |

| Optional Argument | Description |
|---|---|
| -E "*character_set*" | Specifies native character set encoding.<br>**See Also:** "Using NLS with ldifwrite" on page 12-9 |

The following example writes all the entries under `ou=Europe, o=imc, c=us` into the `output1.ldi` file.

```
ldifwrite -c nldap -b "ou=Europe, o=imc, c=us" -f output1.ldi
```

All the arguments are mandatory.

The LDIF file and the intermediate file are always written to the current directory.

The ldifwrite tool includes the operational attributes of each entry in the directory, including `createtimestamp`, `creatorsname`, and `orclguid`.

# Catalog Management Tool Syntax

Oracle Internet Directory uses indexes to make attributes available for searches. When Oracle Internet Directory is installed, the entry `cn=catalogs` lists available attributes that can be used in a search. Only those attributes that have an equality matching rule can be indexed.

If you want to use additional attributes in search filters, you must add them to the catalog entry. You can do this at the time you create the attribute by using Oracle Directory Manager. However, if the attribute already exists, then you can index it only by using the Catalog Management tool.

Before running the Catalog Management tool, unset the LANG variable. After you finish running Catalog Management tool, set the LANG variable back to its original value.

To unset LANG:

- Using Korn shell:

  ```
  UNSET LANG
  ```

- Using C shell:

  ```
  UNSETENV LANG
  ```

The Catalog Management tool uses this syntax:

`catalog.sh -connect` *`net_service_name`* `{add|delete} {-attr` *`attr_name`*`|-file` *`filename}`*

| Mandatory Argument | Description |
|---|---|
| - connect *net_service_name* | Specifies the net service name to connect to the directory database<br>**See Also:** *Net8 Administrator's Guide* |

| Optional Arguments | Descriptions |
|---|---|
| - add -attr *attr_name* | Indexes the specified attribute |
| - delete -attr *attr_name* | Drops the index from the specified attribute |
| - add -file *filename* | Indexes attributes (one per line) in the specified file |
| -delete -file *filename* | Drops the indexes from the attributes in the specified file |

When you enter the `catalog.sh` command, the following message appears:

```
This tool can only be executed if you know the OiD user password.
Enter OiD password:
```

If you enter the correct password, the command is executed. If you give an incorrect password, the following message is displayed:

```
Cannot execute this tool
```

After you finish running the Catalog Management tool, set the LANG variable back to its original value.

To set LANG:

- Using Korn shell:

  ```
  SET LANG=appropriate_language; EXPORT LANG
  ```

- Using C shell:

  ```
  SETENV LANG appropriate_language
  ```

To effect the changes after running the Catalog Management tool, stop, then restart, the Oracle directory server.

> **See Also:** Chapter 5, "Managing an Oracle Directory Server" for instructions on starting and restarting directory servers

# OID Monitor Syntax

This section contains these topics:



## Starting the OID Monitor

To start the OID Monitor:

1. Set the following environment variable to the appropriate language setting. The default language set at installation is AMERICAN_AMERICA.

   `NLS_LANG=`*`APPROPRIATE_LANGUAGE`*`.UTF8`

2. At the system prompt, type:

   `oidmon [connect=`*`net_service_name`*`] [sleep=`*`seconds`*`] start`

| Argument | Description |
|---|---|
| connect=*net_service_name* | Specifies the net service name of the database to which you want to connect. This is the network service name set in the `tnsnames.ora` file. This argument is optional. |
| sleep=*seconds* | Specifes number of seconds after which the OID Monitor should check for new requests from OID Control and for requests to restart any servers that may have stopped. The default sleep time is 10 seconds. This argument is optional. |
| start | Starts the OID Monitor process |

For example:

```
oidmon connect=dbs1 sleep=10 start
```

### Stopping the OID Monitor

To stop the OID Monitor daemon, at the system prompt, type:

```
oidmon [connect=net_service_name] stop
```

| Argument | Description |
|---|---|
| connect=*net_service_name* | Specifies net service name of the database to which you want to connect. This is the net service name set in the `tnsnames.ora` file. |
| stop | Stops the OID Monitor process |

For example:

```
oidmon connect=dbsl stop
```

## OID Control Utility Syntax

**Note:** OID Monitor must be running whenever you start, stop, or restart directory server instances.

This section contains these topics:

- Starting and Stopping an Oracle Directory Server Instance
- Starting and Stopping an Oracle Directory Replication Server Instance
- Restarting Directory Server Instances
- Troubleshooting Directory Server Instance Startup

## Starting and Stopping an Oracle Directory Server Instance

Use the **OID Control Utility** to start and stop Oracle directory server instances.

### Starting an Oracle Directory Server Instance

The syntax for starting an Oracle directory server instance is:

```
oidctl connect=net_service_name server=oidldapd instance=server_instance_number
[configset=configset_number] [flags=' -p port_number -work maximum_number_of_
worker_threads_per_server -debug debug_level -l change-logging -server n'] start
```

| Argument | Description |
|---|---|
| `connect` | If you already have a `tnsnames.ora` file configured, this is the net service name specified in that file, located in `ORACLE_HOME/network/admin` |
| `server` | Type of server to start (valid values are OIDLDAPD and OIDREPLD). This is not case-sensitive. |
| `instance` | Instance number of the server to start. Should be a number between 0 and 1000. |
| `configset` | Configset number used to start the server. This defaults to `configset0` if not set. This should be a number between 0 and 1000. |
| `-p` | Specifies a port number during server instance startup. Default port if not set is 389. |
| `-work` | Specifies the maximum number of worker threads for this server |
| `-debug` | Specifies a debug level during Oracle directory server instance startup |
| `-l` | Turns replication change-logging on and off. To turn it off, enter -l. To turn it on, omit the flag. The default is true (values = true and false). (directory server only) |
| `-server` | Specifies the number of server processes to start on this port |
| `start` | Starts the server specified in the `server` argument. |

For example, to start an Oracle directory server instance whose net service name is dbs1, using configset5, at port 12000, with a debug level of 1024, an instance number 3, and in which change-logging is turned off, type at the system prompt:

```
oidctl connect=dbs1 server=oidldapd instance=3 configset=5 flags='-p 12000
-debug 1024 -l ' start
```

When starting and stopping an Oracle directory server instance, the server name and instance number are mandatory. All other arguments are optional.

All keyword value pairs within the flags arguments must be separated by a single space.

Single quotes are mandatory around the flags.

The configset identifier defaults to zero (`configset0`) if not set.

---

**Note:** If you choose to use a port other than the default port (389 for non-secure usage or 636 for secure usage), you must tell the clients which port to use to locate the Oracle Internet Directory. If you use the default ports, clients can connect to the Oracle Internet Directory without referencing a port in their connect requests.

---

### Stopping an Oracle Directory Server Instance

At the system prompt, type:

```
oidctl connect=net_service_name server=OIDLDAPD instance=server_instance_number
stop
```

For example:

```
oidctl connect=dbs1 server=oidldapd instance=3 stop
```

## Starting and Stopping an Oracle Directory Replication Server Instance

Use the OID Control Utility to start and stop Oracle directory replication server instances.

### Starting an Oracle Directory Replication Server Instance

The syntax for starting the Oracle directory replication server is:

```
oidctl connect=net_service_name server=oidrepld instance=server_instance_number
[configset=configset_number] flags=' -h hostname -p port_number
-d debug_level -z transaction_size ' start
```

| Argument | Description |
|---|---|
| `connect` | If you already have a `tnsnames.ora` file configured, then this is the name specified in that file, which is located in ORACLE_HOME/network/admin |
| `server` | Type of server to start (valid values are OIDLDAPD and OIDREPLD). This is not case-sensitive. |
| `instance` | Instance number of the server to start. Should be a number between 0 and 1000. |
| `configset` | Configset number used to start the server. This defaults to configset0 if not set. This should be a number between 0 and 1000. |
| `-p` | Specifies a port number during server instance startup. Default port if not set is 389. |
| `-d` | Specifies a debug level during replication server instance startup |
| `-h` | Specifies the host name on which the server runs. (Replication server only) |
| `-m [true\|false]` | Turns conflict resolution on and off. The default is true (values = true and false). (Replication server only) |
| `-z` | Specifies the number of changes applied in each replication update cycle. If you do not specify this, the number is determined by the Oracle directory server sizelimit parameter, which has a default setting of 1024. You can configure this latter setting. |
| `start` | Starts the server specified in the *server* argument. |

For example, to start the replication server with an instance=1, at port 12000, with debugging set to 1024, type at the system prompt:

```
oidctl connect=dbs1 server=oidrepld instance=1 flags='-p 12000 -h eastsun11 -d
1024' start
```

When starting and stopping an Oracle directory replication server, the `-h` flag, which specifies the host name, is mandatory. All other flags are optional.

All keyword value pairs within the flags arguments must be separated by a single space.

Single quotes are mandatory around the flags.

The configset identifier defaults to zero (`configset0`) if not set.

> **Note:** If you choose to use a port other than the default port (389 for non-secure usage or 636 for secure usage), you must tell the clients which port to use to locate the Oracle Internet Directory. If you use the default ports, clients can connect to the Oracle Internet Directory without referencing a port in their connect requests.

### Stopping an Oracle Directory Replication Server Instance

At the system prompt, type:

```
oidctl connect=net_service_name server=OIDREPLD instance=server_instance_number
stop
```

For example:

```
oidctl connect=dbs1 server=oidrepld instance=1 stop
```

## Restarting Directory Server Instances

To restart a directory server instance, at the system prompt, type:

```
oidctl connect=net_service_name server={oidldapd|oidrepld} instance=server_
instance_number  restart
```

OID Monitor must be running whenever you start, stop, or restart directory server instances.

If you try to contact a server that is down, you receive from the SDK the error message `81—LDAP_SERVER_DOWN`.

If you change a configuration set entry that is referenced by an active server instance, you must stop that instance and restart it to effect the changed value in the configuration set entry on that server instance. You can either issue the STOP command followed by the START command, or you can use the RESTART command. RESTART both stops and restarts the server instance.

For example, suppose that Oracle directory server instance1 is started, using configset3, and with the net service name dbs1. Further, suppose that, while instance1 is running, you change one of the attributes in configset3. To enable the change in configset3 to take effect on instance1, you enter the following command:

```
oidctl connect=dbs1 server=oidldapd instance=1 restart
```

If there are more than one instance of the Oracle directory server running on that node using configset3, then you can restart all the instances at once by using the following command syntax:

```
oidctl connect=dbs1 server=oidldapd restart
```

Note that this command restarts all the instances running on the node, whether they are using configset3 or not.

**Important Note:** During the restart process, clients cannot access the Oracle directory server instance. However, the process takes only a few seconds to execute.

## Troubleshooting Directory Server Instance Startup

If the directory server fails to start, you can override all user-specified configuration parameters to start the directory server and then return the configuration sets to a workable state by using the ldapmodify operation.

To start the directory server using its hard-coded default parameters instead of the configuration parameters stored in the directory, type at the system prompt:

```
oidctl connect=net_service_name flags='-p port_number -f'
```

The `-f` option in the flags starts the server with hard-coded configuration values, overriding any defined configuration sets except for the values in `configset0`.

# OID Database Password Utility Syntax

The OID Database Password Utility syntax is:

```
oidpasswd [connect=net_service_name]
```

The OID Database Password Utility prompts you for the current password. Type the current password, then the new password, then a confirmation of the new password.

The OID Database Password Utility assumes by default that the password being changed is that of the local database (as defined by *ORACLE_HOME* and ORACLE_*SID*). If you are changing the password on a remote database, you must use the `connect=net_service_name` option.

For example:

```
$ oidpasswd
current password: ods
new password: newsupersecret
confirm password: newsupersecret
password set.$
```

**Note:** User responses are not echoed to the screen.

# OID Database Statistics Collection Tool Syntax

The $*ORACLE_HOME*/ldap/admin/oidstats.sh tool is provided to analyze the various database `ods` schema objects to estimate the statistics.

## Syntax

```
oidstats.sh [ -connect database_connect_string ]
            [ -login database_account_login ]
            [ -pass database_account_password ]
            [ -all ]
            [ -cat catalog_name ]
            [ -pct percent ]
            [ -help | -usage ]
```

## Parameters

| Parameter | Description | Default |
|---|---|---|
| connect | DB connect string | *ORACLE_SID* |
| login | DB user name | `ods` |
| pass | DB Account Password | `ods` |
| all | Estimate statistics on all catalog tables plus DN catalogue | All catalogues |
| cat | Estimate statistics either on all catalogs (all) or on a particular one, for example, ct_cn | None |
| pct | Percent of Data To sample | 100 |

## Examples: Using the OID Database Statistics Collection Tool

Each of the following examples assume that the ORACLE_*SID* and the default user name and password are in effect.

This example estimates statistics based on 100 percent sample data of all tables.

```
oidstats.sh -all -pct 100
```

This example estimates statistics based on 50 percent sample data of all tables.

```
oidstats.sh -all -pct 50
```

This example estimates statistics based on 50 percent sample data of CT_CN table.

```
oidstats.sh -cat ct_cn -pct 50
```

This example estimates statistics based on 40 percent sample data of all catalog tables

```
oidstats.sh -cat all -pct 40
```

# B

# Adding a DSA Using the Database Copy Procedure

This appendix describes the procedure to add a new **DSA** to an existing replicating system by using the database copy procedure, also known as **cold backup**.

> **Note:** Because this procedure involves copying Oracle data files, faster performance depends on the underlying network. If the underlying network is weak, then it may be better to implement the method described in Chapter 10, "Managing Directory Replication", or to physically ship compressed Oracle data files on a medium such as a tape or disk. Consult your local system or network administrator for more details on the network.
>
> Only a person familiar with the Oracle database should implement this procedure.

This appendix contains these topics:

## Assumptions

This document assumes that the UNIX directories are created according to Optimal Flexible Architecture (OFA), the set of configuration guidelines for efficient and reliable Oracle databases.

> **See Also:** The Oracle installation guide for your operating system for more information on OFA

## Sponsor Directory Site Environment

Set up the environment of the sponsor site. In the example shown throughout this chapter, the host name is rst-sun.

```
Hostname    = rst-sun
ORACLE_BASE = /private/oracle/app/oracle
ORACLE_HOME = /private/oracle/app/oracle/product/8.1.6
ORACLE_SID  = LDAP
LD_LIBRARY_PATH = $ORACLE_HOME/lib
NLS_LANG    = AMERICAN_AMERICA.UTF8
datafile location = /private/oracle/oradata/LDAP
Dump destination =  /private1/oracle/app/oracle/admin/LDAP/pfile,
                    /private1/oracle/app/oracle/admin/LDAP/bdump,
                    /private1/oracle/app/oracle/admin/LDAP/cdump,
                    /private1/oracle/app/oracle/admin/LDAP/udump,
                    /private1/oracle/app/oracle/admin/LDAP/create
```

## New Directory Site Environment

Set up the environment for the new directory site. In the example shown throughout this chapter, the new site is on the node named dsm-sun.

```
Hostname = dsm-sun
ORACLE_BASE = /private1/oracle/app/oracle
ORACLE_HOME = /private1/oracle/app/oracle/product/8.1.6
ORACLE_SID  = NLDAP
LD_LIBRARY_PATH = $ORACLE_HOME/lib
NLS_LANG = AMERICAN_AMERICA.UTF8
   datafile location = /private1/oracle/oradata/NLDAP
   Dump destination =  /private1/oracle/app/oracle/admin/NLDAP/pfile,
                       /private1/oracle/app/oracle/admin/NLDAP/bdump,
                       /private1/oracle/app/oracle/admin/NLDAP/cdump,
                       /private1/oracle/app/oracle/admin/NLDAP/udump,
                       /private1/oracle/app/oracle/admin/NLDAP/create
```

---

**Note:** After installation of the Oracle database or Oracle directory, you use Oracle Database Configuration Assistant to create data file directories. Create the new directories on the new node under various UNIX partitions as defined by OFA.

---

## Tasks To Be Performed on the Sponsor Node

Complete the following steps on the sponsor node.

1. At the command line prompt execute SQL*Plus.

```
$ sqlplus
SQL> connect internal
SQL> ALTER DATABASE BACKUP CONTROLFILE TO TRACE;
```

The above command will create a trace file under the user dump destination directory (that is, `/private1/oracle/app/oracle/admin/LDAP/udump`).

The file will be created in the following format:

`$ORACLE_SID_<ora_processid>.trc`

For example:

`ldap_ora_4765.trc`

2. Shutdown the LDAP and replication servers and OID Monitor processes. Make sure the ldap and replication servers are stopped before stopping the OID Monitor process.

```
$ oidctl connect=<net_service_name> server=oidrepld instance=<inst_#> stop
$ oidctl connect=<net_service_name> server=oidldapd instance=<inst_#> stop
$ oidmon connect=<net_service_name> stop
```

   In these commands, *net_service_name* is the net service name in the node's `tnsnames.ora` file.

3. On the remaining nodes, shutdown the LDAP replication server only.

```
$ oidctl connect=<net_service_name> server=oidrepld instance=<inst_#> stop
```

   Repeat the above procedure on all nodes except the sponsor node. Specify appropriate net service names for the corresponding nodes.

4. Quiesce **Advanced Symmetric Replication (ASR)** by running the following script at the **master definition site (MDS)**:

```
ldaprepl.sh -quiesce
```

   Enter the Oracle global name for the MDS when prompted.

   **Note:** This procedure can take place only on the Master Definition Site.

   At this point, other nodes are available for LDAP edits only, but replication will not take place.

5. After quiescing the environment, shutdown the database and Net8 listener on the sponsor node only:

```
$ lsnrctl [listener name] stop   (By default listener name is LISTENER)
$ sqlplus
SQL> connect internal
SQL> shutdown normal
SQL> exit
```

6. Copy the trace file created under Step 1 to a new file, `newdb.sql`, under the same directory.

```
$ cd $ORACLE_BASE/admin/LDAP/udump
$ cp ldap_ora_4765.trc newdb.sql
```

7. Edit `newdb.sql`, using any text editor, and delete the lines up to START NOMOUNT.

```
CREATE CONTROLFILE REUSE SET DATABASE <database_name> RESETLOG
```

8. Modify the UNIX directory location of the database/logfiles etc. to point to the new node directory. Refer to the sample file `newdb.sql` as follows:

```
Begin newdb.sql
CREATE CONTROLFILE REUSE SET DATABASE "LDAP" RESETLOGS
MAXLOGFILES 16
MAXLOGMEMBERS 2
MAXDATAFILES 255
MAXINSTANCES 1
MAXLOGHISTORY 100
LOGFILE
GROUP 1 '/private2/oracle/oradata/NLDAP1/log1_NLDAP.dbf'  SIZE 1M,
GROUP 2 '/private2/oracle/oradata/NLDAP1/log2_NLDAP.dbf'  SIZE 1M
DATAFILE
'/private2/oracle/oradata/NLDAP1/sys0_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/rbs1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/attrs1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/dncat1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/cncat1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/objcl1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/cats1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/default1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/temp1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/iattrs1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/idncat1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/icncat1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/iobjcl1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/icats1_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/temp2_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/cats2_NLDAP.dbf',
'/private2/oracle/oradata/NLDAP1/attrs2_NLDAP.dbf'
;
 End newdb.sql
```

**9.** Copy the files `initLDAP.ora` and `configLDAP.ora` under `$ORACLE_HOME/dbs` to `initNLDAP.ora` and `configNLDAP.ora` respectively.

```
$cd $ORACLE_HOME/dbs
$cp initLDAP.ora initNLDAP.ora
$cp configLDAP.ora configNLDAP.ora
```

**10.** Edit the copied file (`initNLDAP.ora`) and comment out the parameter JOB_QUEUE_PROCESS. Change the following parameter:

db_name = LDAP (If the parameter does not exist in the file initNLDAP.ora, then modify the file configNLDAP.ora)
ifile = *UNIX_directory_location_of_the_new_config_file/* configNLDAP.ora

**11.** Edit the copied file `configNLDAP.ora` to change the following parameters:

cdump = *UNIX_directory_location_of_the_new_node*
udump = *UNIX_directory_location_of_the_new_node*
bdump = *UNIX_directory_location_of_the_new_node*
control_files = *UNIX_directory_location_of_the_new_node*

**12.** Edit the `tnsnames.ora` file to include information pertaining to the new node. Refer to the following sample file:

```
Begin tnsnames.ora

ldap1.world =
   (description=
      (address=(protocol=tcp)(host=rst-sun)(port=1521))
      (connect_data=(sid=LDAP))
   )
ldap2.world =
   (description=
      (address=(protocol=tcp)(host=eas-sun10)(port=1521))
      (connect_data=(sid=LDAP))
   )
ldap3.world =
   (description=
      (address=(protocol=tcp)(host=dsm-sun)(port=1521))
      (connect_data=(sid=NLDAP))
   )

End tnsnames.ora
```

**13.** Copy the file `listener.ora` to `list.bak`. Edit the copied file `list.bak` to include the information pertaining to the new node. Refer to the following sample file:

```
Begin listener.ora

# The KEY value for the IPC protocol may be anything, and
# is not related to either the TCP hostname or database SID.

LISTENER =
  (ADDRESS_LIST =
        (ADDRESS= (PROTOCOL= IPC)(KEY= LDAP))
        (ADDRESS= (PROTOCOL= IPC)(KEY= PNPKEY))
        (ADDRESS= (PROTOCOL= TCP)(Host= dsm-sun)(Port= 1521))
  )
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME= dsm-sun.us.oracle.com)
      (ORACLE_HOME= /private1/oracle/app/oracle/product/8.1.6)
      (SID_NAME = NLDAP)
    )
    (SID_DESC =
      (SID_NAME = extproc)
      (ORACLE_HOME = /private1/oracle/app/oracle/product/8.1.6)
      (PROGRAM = extproc)
    )
  )
STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10
TRACE_LEVEL_LISTENER = OFF

End listener.ora
```

The files `tnsnames.ora` and `listener.ora` can reside under `$ORACLE_HOME/network/admin` or `/var/opt/oracle` or under the directory pointed to by the TNS_ADMIN environment variable.

**14.** Copy the updated `tnsnames.ora` file to all the nodes. Be careful to copy it to the location of the current `tnsnames.ora` on each node. The file `tnsnames.ora` can be copied to other nodes using FTP. Make sure you transfer the file in ASCII mode.

Prior to copying the file `tnsnames.ora` to the new node, install the Oracle database software on the new node. Also copy the files `list.bak` as `listener.ora` and `sqlnet.ora` from the sponsor node to the new node.

**15.** Create an archive of all the data files and compress the archived file. For example:

```
$ >oradb.tar
```

This command will create an empty file under a directory. Make sure you have enough space in the partition where the archives will be created.

```
$ find / -name *.dbf -print -exec tar rvf  <absolute_path_of_the_directory_
which_contains_oradb.tar> {} \;
```

This command will search for all files ending with extension `.dbf` from the root directory. The assumption is that there is only one instance of the database server installed on the node and data files end with `*.dbf` extension.

```
$ find / -name *.log -print -exec tar rvf <absolute_path_of_the_directory_
which_contains_oradb.tar>
$ compress oradb.tar
```

This procedure is only an example to illustrate the method to back up the files. The Oracle data files will be backed up in the absolute path using this method. It is a better idea to back up the files from the current directory, so that you have more flexibility when you want to restore the data files. Consult your system administrator before backing up the database.

## Tasks To Be Performed on the New Node

Complete the following steps on the new node.

1. Log in to the new node (dsm-sun).
2. Edit the `oratab` file appropriately for the new instance, at all database nodes. See the sample file for syntax.

   ```
   Begin oratab

   NLDAP:/private1/oracle/app/oracle/product/8.1.6:N
   *:/private1/oracle/app/oracle/product/8.1.6:N

   End oratab
   ```

3. Make sure the environment variables are set in the new directory site.
4. Install the Oracle database and Oracle directory server. Perform software only install of the Oracle database and directory server. Installation of Oracle database and directory software can be performed on the new node at any time before the database files are copied to the new machine. Perform post-installation (that is: `root.sh`) activities for the database as well as the Directory server.

   **See Also:** Oracle8*i* installation

   If you have already performed Oracle database and Directory installation on the new node, then proceed to Step 5.
5. Copy the files `initNLDAP.ora` and `configNLDAP.ora` from the sponsor node (rst-sun) to the new node under the UNIX directory `$ORACLE_BASE/ADMIN/NLDAP/PFILE`. Files can be copied to the new machine using tools such as FTP. Make sure the transfer mode is ASCII.
6. Create a symbolic soft link from `$ORACLE_HOME/DBS TO $ORACLE_BASE/ADMIN/NLDAP/PFILE`.

   ```
   $ ln -s $ORACLE_BASE/admin/NLDAP/pfile/initNLDAP.ora
          $ORACLE_HOME/dbs/initNLDAP.ora
   $ ln -s $ORACLE_BASE/admin/NLDAP/pfile/configNLDAP.ora
          $ORACLE_HOME/dbs/configNLDAP.ora
   ```

7. Copy the archived file created in the sponsor node procedure, using a tool such as FTP. (You created this file in Step 15 on page B-8.) Set the transfer mode to binary.

```
ftp> open rst-sun
Connected to rst-sun.us.oracle.com.
220 rst-sun FTP server (UNIX(r) System V Release 4.0) ready.
Name (rst-sun:oracle):
331 Password required for oracle.
Password:
230 User oracle logged in.
ftp> cd /private1/oracle/oradata/LDAP
250 CWD command successful.
ftp> binary
200 Type set to I.
ftp> mget oradb.tar.Z
```

   If the data files are huge (several gigabytes or terabytes) and the network bandwidth is low, then it may be a better idea to physically ship the compressed file on any media, such as tape or disk, from the sponsor to the new node.

8. Copy the file `newdb.sql` created under Step 6 of the sponsor node setup to the background user dump destination directory. You must transfer the file newdb.sql only in ASCII mode. For example:

```
 $ cd /private1/oracle/app/oracle/admin/NLDAP/udump
                    (that is::$ORACLE_BASE/admin/<SID>/udump)
$ ftp
ftp> open rst-sun
ftp> cd /private1/oracle/app/oracle/admin/LDAP/udump
ftp> mget newdb.sql
```

9. At the UNIX shell prompt execute the following commands:

```
$ sqlplus
SQL> connect internal
SQL> startup nomount
SQL> @newdb.sql
SQL> shutdown normal
SQL> startup (uncomment the parameter job_queue_process prior to  startup)
SQL>exit
$ lsnrctl start
```

**10.** Log in to the sponsor node and start up the database and listener on the sponsor node; for example, rst-sun.

```
$ telnet rst-sun
$ sqlplus
SQL> connect internal
SQL> startup
SQL> exit
$ lsnrctl start (By default listener name is LISTENER)
$ exit
```

**11.** If the sponsor node is a master site, then proceed to Step 12.

If the new node is created by using backup database copy of the MDS, then the master definition catalog needs to be dropped and the underlying ASR catalogs must be created. To drop the definition of the MDS from the ASR catalog on the new node and add the ASR catalogs, execute the following scripts.

```
$ cd $ORACLE_HOME/ldap/admin
$ sqlplus repadmin/repadmin
SQL> @ldapdropmds.sql
SQL> @ldapcreindex.sql
```

Specify the global name of the new node when prompted.

**12.** To configure the ASR, at the shell prompt, execute the following command:

```
$ ldaprepl.sh -addnode
```

**13.** Update the LDAP replication agreements to include the new node.

Sample LDIF file:

```
dn: orclagreementid=000001, cn=orclreplagreements
changetype: modify
add: orcldirreplgroupdsas
orcldirreplgroupdsas: dsm-sun
```

**14.** Start up the LDAP replication server on all the nodes, including new and sponsor nodes.

# Verification Process

Log in to the Oracle database by using SQL*Plus and specify the user name as ODS, and the password ods when prompted.

Check the ods_chg_stat table on all nodes and see if they have correct and identical rows. The ods_chg_stat table should contain (*number of nodes*) x (*number of nodes*) rows. For example, if there were two nodes participating in ASR-based replication, and you added a third node, the ods_chg_stat table would contain nine rows, that is, 3 x 3, on each node. The rows are shown in the following table:

| Supplier | Consumer | Change Number |
|---|---|---|
| Node1 | node2 | <number 1> |
| Node1 | node3 | <number 2> |
| Node1 | node1 | <number 3> |
| Node2 | node1 | <number 4> |
| Node2 | node2 | <number 5> |
| Node2 | node2 | <number 6> |
| Node3 | node1 | 0 |
| Node3 | node2 | 0 |
| Node3 | node3 | 0 |

The rows with consumer names identical to that of suppliers contain the last changes processed by the outbound change log processing threads at the supplier sides. The rows with different supplier and consumer names contain last change numbers already processed from the suppliers to the consumers in question.

Since Node3 is a new node, there have been no changes supplied by Node3 yet. Therefore, the change numbers for Node3 as supplier are 0.

There may be a time delay before all nodes contain identical rows, but this delay should not be more than two to three minutes.

# C

# Using Oracle Wallet Manager

Security administrators use Oracle Wallet Manager to manage public-key security credentials on Oracle clients and servers. The wallets it creates are opened by using either the Oracle Enterprise Login Assistant or the Oracle Wallet Manager.

This chapter describes the Oracle Wallet Manager, in the following sections:

- Overview
- Managing Wallets
- Managing Certificates

> **See Also:** *Oracle Advanced Security Administrator's Guide* for information about how to open and close wallets for secure SSL communications by using Oracle Enterprise Login Assistant

# Overview

Traditional private-key or symmetric-key cryptography requires that entities desiring to establish secure communications possess a single secret key known only to them. *Harriet* and *Dick*, for example, could agree to shift each letter in their private messages by two character positions (A becomes C, B becomes E, and so on) to encrypt the message text. Using this method, a *HELLO* message from Harriet to Dick would read *JGNNP*. The actual encryption methods in current use are much more complex and significantly more secure, but an underlying problem remains—sending messages encrypted with a single key requires prior, *secure* distribution of the key to each participating party. Otherwise, a malicious third party might obtain the key, intercept communications, and compromise security. Public-key cryptography addresses this problem, by providing a secure method for key distribution.

Public-key cryptography requires a party to possess a **public/private key pair**. The **private key** is kept secret and is known only to that party. The **public key**, as the name implies, is freely available. To send a secret message to this party requires that a third party sender encrypt the message with the public key. Such a message can only be decrypted by a party holding the associated private key.

For example, when Dick wants to send a secure message to Harriet, he first asks Harriet for her public key (or obtains it from another, public source). Harriet gives Dick the public key, but Tom, a malicious eavesdropper, also obtains the public key. Nevertheless, when Dick sends Harriet a message encrypted with her public key, Tom cannot decrypt it; the message can only be decrypted with Harriet's private key.

Public-key algorithms thus guarantee the secrecy of a message, but they don't guarantee *secure communications* because they don't verify the identities of the communicating parties. In order to establish secure communications, it is important to verify that the public key used to encrypt a message does in fact belong to the target recipient. Otherwise, a third party can potentially eavesdrop on the communication and intercept public key requests, substituting its public key for a legitimate key.

If Tom, for example, is able to substitute his public key for Harriet's public key and send it to Dick, Dick might then send a message to Harriet encrypted with Tom's public key—believing he was using Harriet's public key. Tom could then decrypt a subsequent intercepted message from Dick using his private key, re-encrypt it with Harriet's public key and re-transmit it to Harriet. Harriet could then decrypt the incoming message using her private key, and never know that it had been intercepted by Tom.

In order to avoid such a man-in-the-middle attack, it is necessary to verify the owner of the public key, a process called **authentication**. This authentication can be accomplished through a **certificate authority (CA)**.

A CA is a third party that is trusted by both of the parties attempting secure communication. The CA issues public key certificates that contain an entity's name, public key, and certain other security credentials. Such credentials typically include the CA name, the CA signature, and the certificate effective dates (From Date, To Date).

The CA uses its private key to encrypt a message, while the public key is used to decrypt it, thus verifying that the message was encrypted by the CA. The CA public key is well known, and does not have to be authenticated each time it is accessed. Such CA public keys are stored in a **wallet**.

Oracle Wallet Manager is a stand-alone Java application that wallet owners use to manage and edit the security credentials in their Oracle wallets. These tasks include the following:

- Generating a public/private key pair and creating a certificate request for submission to a CA.
- Installing a certificate for the entity.
- Configuring **trusted certificates** for the entity.
- Opening a wallet to enable access to PKI-based services.
- Creating a wallet that can be accessed by using either Oracle Enterprise Login Assistant or Oracle Wallet Manager.

# Managing Wallets

This section describes how to create a new wallet and perform associated wallet management tasks, such as generating certificate requests, exporting certificate requests, and importing certificates into wallets, in the following subsections:



## Starting Oracle Wallet Manager

To start Oracle Wallet Manager:

| | |
|---|---|
| UNIX: | Enter `owm` at the command line. |
| Windows NT: | Press Start > *ORACLE_HOME* > Network Administration > Wallet Manager |

## Creating a New Wallet

Create a new wallet as follows:

1. Choose `Wallet > New` from the menu bar; the New Wallet dialog box appears.
2. Read the recommended guidelines for creating a password and enter a password in the Wallet Password field.

   Because an Oracle wallet contains a user's credentials that can be used to authenticate the user to multiple databases, it is especially important to choose

a strong password for the wallet. A malicious user who guesses the password to a user's wallet can access all the databases that the user can access.

Oracle Corporation recommends that you choose a password that is not too short, not easily guessed, and is reasonably complex. A reasonably complex password has at least six characters, and contains at least one symbol or number—so that it will not be found in a dictionary.

Example: `gol8fer`

It is also a prudent security practice for users to change their passwords periodically, such as once a month, or once a quarter.

3. Re-enter that password in the Confirm Password field.
4. Choose `OK` to continue.
5. An Alert is displayed, and informs you that a new empty wallet has been created. It prompts you to decide whether you want to create a certificate request. See: "Creating a Certificate Request" on page C-9.

   If you choose Cancel, you are returned to the Oracle Wallet Manager main window. The new wallet you just created appears in the left window pane. The certificate has a status of `Empty`, and the wallet displays its default trusted certificates.
6. Select `Wallet > Save In System Default` to save the new wallet.

   If you do not have permission to save the wallet in the system default, you can save it to another location.

   A message at the bottom of the window informs you that the wallet was successfully saved.

## Opening an Existing Wallet

Open a wallet that already exists in the file system directory as follows:

1. Choose `Wallet > Open` from the menu bar; the Select Directory dialog box appears.
2. Navigate to the directory location in which the wallet is located, and select the directory.
3. Choose OK; the Open Wallet dialog box appears.
4. Enter the wallet password in the Wallet Password field.

5. Choose OK.
6. The message `Wallet opened successfully` appears at the bottom of the window, and you are returned to the Oracle Wallet Manager main window. The wallet's certificate and its trusted certificates are displayed in the left window pane.

## Closing a Wallet

To close an open wallet in the currently selected directory:

- Choose `Wallet > Close`.
- The message `Wallet closed successfully` appears at the bottom of the window, to confirm that the wallet is closed.

## Saving Changes

To save your changes to the current open wallet:

- Choose `Wallet > Save`.
- A message at the bottom of the window confirms that the wallet changes were successfully saved to the wallet in the selected directory location.

## Saving the Open Wallet to a New Location

Use the `Save As` option to save the current open wallet to a new directory location:

1. Choose `Wallet > Save As`. The select directory dialog box appears.
2. Select a directory location to save the wallet.
3. Choose OK.

   The following message appears if a wallet already exists in the selected directory:

   ```
   A wallet already exists in the selected path. Do you
   want to overwrite it?.
   ```

   Choose Yes to overwrite the existing wallet, or No to save the wallet to another directory.

   A message at the bottom of the window confirms that the wallet was successfully saved to the selected directory location.

## Saving in System Default

Use the `Save in System Default` menu option to save the current open wallet to the system default directory location. This makes the current open wallet the wallet that is used by SSL:

- Choose `Wallet > Save in System Default`.
- A message at the bottom of the window confirms that the wallet was successfully saved in the system default wallet location.

## Deleting the Wallet

To delete the current open wallet:

1. Choose `Wallet > Delete`; the `Delete Wallet` dialog box appears.
2. Review the displayed wallet location to verify you are deleting the correct wallet.
3. Enter the wallet password.
4. Choose `OK`; a dialog panel appears to inform you that the wallet was successfully deleted.

---

**Note:** Any open wallet in application memory will remain in memory until the application exits. Therefore, deleting a wallet that is currently in use does not immediately affect system operation.

---

## Changing the Password

A password change is effective immediately. The wallet is saved to the currently selected directory, with the new encrypted password.To change the password for the current open wallet:

1. Choose `Wallet > Change Password;` the `Change Wallet Password` dialog box appears.
2. Enter the existing wallet password.
3. Enter the new password.
4. Re-enter the new password.
5. Choose `OK`.

A message at the bottom of the window confirms that the password was successfully changed.

## Using Auto Login

The Oracle Wallet Manager Auto Login feature opens a copy of the wallet and enables PKI-based access to secure services—as long as the wallet in the specified directory remains open in memory.

You must enable Auto Login if you want single sign-on access to multiple Oracle databases.

### Enabling Auto Login

To enable Auto Login:

1. Choose `Wallet` from the menu bar.
2. Choose the check box next to the Auto Login menu item; a message at the bottom of the window displays `Autologin enabled`.

### Disabling Auto Login

To disable Auto Login:

1. Choose `Wallet` from the menu bar.
2. Choose the check box next to the Auto Login menu item; a message at the bottom of the window displays `Autologin disabled`.

## Using Oracle Wallet Manager with Oracle Application Server

When using the Oracle Application Server (OAS), you must install the Oracle Wallet Manager on a primary node and on each remote node in a multi-node configuration. After you install the product on each node you must then copy the wallet from the primary node to each of the remote nodes.

# Managing Certificates

Oracle Wallet Manager uses two kinds of certificates: user certificates and trusted certificates. This section describes how to manage both certificate types, in the following subsections:

- Managing User Certificates
- Managing Trusted Certificates

> **Note:** You must first install a trusted certificate from the certificate authority before you can install a user certificate issued by that authority. Several trusted certificates are installed by default when you create a new wallet.

## Managing User Certificates

Managing user certificates involves the following tasks:

- Creating a Certificate Request
- Exporting a User Certificate Request
- Importing the User Certificate into the Wallet
- Removing a User Certificate from a Wallet

### Creating a Certificate Request

The actual certificate request becomes part of the wallet. You can reuse any certificate request to obtain a new certificate. However, you cannot edit an existing certificate request; store only a correctly filled out certificate request in a wallet.

To create a PKCS #10 certificate request:

1. Choose `Operations > Create Certificate Request`; the `Create Certificate Request` dialog box appears.
2. Enter the following information (Table C–1):

***Table C–1 Certificate Request: Fields and Descriptions***

| Field Name | Description |
|---|---|
| Common Name | Mandatory. Enter the name of the user's or service's identity. Enter a user's name in first name / last name format. |

***Table C–1 Certificate Request: Fields and Descriptions***

| Field Name | Description |
|---|---|
| Organizational Unit | Optional. Enter the name of the identity's organizational unit. Example: Finance. |
| Organization | Optional.Enter the name of the identity's organization. Example: XYZ Corp. |
| Locality / City | Optional. Enter the name of the locality or city in which the identity resides. |
| State / Province | Optional. Enter the full name of the state or province in which the identity resides.<br>Enter the full state name, because some certificate authorities do not accept two–letter abbreviations. |
| Country | Mandatory. Choose the drop-down list to view a list of country abbreviations. Select the country in which the organization is located. |
| Key Size | Mandatory. Choose the drop-down box to view a list of key sizes to use when creating the public / private key pair. |
| Advanced | Optional. Choose Advanced to view the Advanced Certificate Request dialog panel. Use this field to edit or customize the identity's distinguished name (DN). For example, you can edit the full state name and locality. |

3. Choose OK. An Oracle Wallet Manager dialog box informs you that a certificate request was successfully created. You can either copy the certificate request text from the body of this dialog panel and paste it into an e-mail message to send to a certificate authority, or you can export the certificate request to a file.
4. Choose OK. You are returned to the Oracle Wallet Manager main window; the status of the certificate is changed to Requested.

### Exporting a User Certificate Request

Save the certificate request in a file system directory when you elect to export a certificate request:

1. Choose `Operations > Export Certificate Request` from the menu bar; the Export Certificate Request dialog box appears.
2. Enter the file system directory in which you want o save your certificate request, or navigate to the directory structure under Folders.
3. Enter a file name to save your certificate request, in the Enter File Name field.
4. Choose `OK`. A message at the bottom of the window confirms that the certificate request was successfully exported to the file. You are returned to the Oracle Wallet Manager main window.

### Importing the User Certificate into the Wallet

You will receive an e-mail notification from the certificate authority informing you that your certificate request has been fulfilled. Import the certificate into a wallet in either of two ways: copy and paste the certificate from the e-mail you receive from the certificate authority, or import the user certificate from a file.

#### Pasting the Certificate

To paste the certificate:

1. Copy the certificate text from the e-mail message or file you receive from the certificate authority. Include the lines `Begin Certificate` and `End Certificate.`
2. Choose `Operations > Import User Certificate` from the menu bar; the Import Certificate dialog box appears.
3. Choose the `Paste the Certificate` button, and choose `OK`; an Import Certificate dialog box appears with the following message:

   ```
   Please provide a base64 format certificate and paste it below.
   ```

4. Paste the certificate into the dialog box, and choose `OK`. A message at the bottom of the window confirms that the certificate was successfully installed. You are returned to the Oracle Wallet Manager main panel, and the wallet status changes to `Ready`.

#### Selecting a File that Contains the Certificate

To select the file:

1. Choose `Operations > Import User Certificate` from the menu bar.
2. Choose the `Select a file...` certificate button, and choose `OK`; the Import Certificate dialog box appears.
3. Enter the path or folder name of the certificate location.
4. Select the name of the certificate file (for example, `cert.txt`).
5. Choose `OK`. A message at the bottom of the window appears, to inform you that the certificate was successfully installed. You are returned to the Oracle Wallet Manager main panel, and the wallet status is changes to `Ready`.

### Removing a User Certificate from a Wallet

1. Choose `Operations > Remove User Certificate`; a dialog panel appears and prompts you to verify that you want to remove the user certificate from the wallet.
2. Choose `Yes`; you are returned to the Oracle Wallet Manager main panel, and the certificate displays a status of `Requested`.

## Managing Trusted Certificates

Managing trusted certificates includes the following tasks:

- Importing a Trusted Certificate
- Removing a Trusted Certificate
- Exporting a Trusted Certificate
- Exporting All Trusted Certificates
- Exporting a Wallet

### Importing a Trusted Certificate

You can import a trusted certificate into a wallet in either of two ways: paste the trusted certificate from an e-mail that you receive from the certificate authority, or import the trusted certificate from a file.

Oracle Wallet Manager automatically installs trusted certificates from VeriSign, RSA, and GTE CyberTrust when you create a new wallet.

**Pasting the Trusted Certificate** To paste the trusted certificate:

1. Choose `Operations > Import Trusted Certificate` from the menu bar; the Import Trusted Certificate dialog panel appears.
2. Choose the `Paste the Certificate` button, and choose `OK`. An Import Trusted Certificate dialog panel appears with the following message:

   ```
   Please provide a base64 format certificate and paste it below.
   ```

3. Copy the trusted certificate from the body of the e-mail message you received that contained the user certificate. Include the lines `Begin Certificate` and `End Certificate.`
4. Paste the certificate into the window, and Choose `OK`. A message at the bottom of the window informs you that the trusted certificate was successfully installed.
5. Choose `OK`; you are returned to the Oracle Wallet Manager main panel, and the trusted certificate appears at the bottom of the Trusted Certificates tree.

#### Selecting a File that Contains the Trusted Certificate

To select the file:

1. Choose `Operations > Import Trusted Certificate` from the menu bar. The Import Trusted Certificate dialog panel appears.
2. Enter the path or folder name of the trusted certificate location.
3. Select the name of the trusted certificate file (for example, `cert.txt`).
4. Choose `OK`. A message at the bottom of the window informs you that the trusted certificate was successfully imported into the wallet.
5. Choose `OK` to exit the dialog panel; you are returned to the Oracle Wallet Manager main panel, and the trusted certificate appears at the bottom of the Trusted Certificates tree.

### Removing a Trusted Certificate

To remove a trusted certificate from a wallet:

1. Select the trusted certificate listed in the Trusted Certificates tree.
2. Choose `Operations > Remove Trusted Certificate` from the menu bar.

A dialog panel warns you that your user certificate will no longer be verifiable by its recipients if you remove the trusted certificate that was used to sign it.

3. Choose `Yes`; the selected trusted certificate is removed from the Trusted Certificates tree.

---

**Note:** A certificate that is signed by a trusted certificate is no longer verifiable when you remove it from your wallet.

Also, you cannot remove a trusted certificate if it has been used to sign a user certificate that is still present in the wallet. To remove such a trusted certificate, you must first remove the certificates that it has signed.

---

### Exporting a Trusted Certificate

To export a trusted certificate to another file system location:

1. Select `Operations > Export Trusted Certificate`; the Export Trusted Certificate dialog box appears.
2. Select a file system directory to save your trusted certificate, or choose `Browse` to display the directory structure.
3. Enter a file name to save your trusted certificate.
4. Choose `OK`; you are returned to the Oracle Wallet Manager main window.

### Exporting All Trusted Certificates

To export all of your trusted certificates to another file system location:

1. Choose `Operations > Export All Trusted Certificates`. The Export Trusted Certificate dialog box appears.
2. Select the file system directory to save your trusted certificates, or choose `Browse` to display the directory structure.
3. Enter a file name to save your trusted certificates.
4. Choose `OK`; you are returned to the Oracle Wallet Manager main window.

## Exporting a Wallet

You can export a wallet to text-based PKI formats. Individual components are formatted according to the following standards (Table C–2):

***Table C–2 PKI Wallet Encoding Standards***

| Component | Encoding Standard |
|---|---|
| Certificate chains | X509v3 |
| Trusted certificates | X509v3 |
| Private keys | PKCS5 |

# D

# Using Access Control Directive Format

This appendix describes the format (syntax) of any **Access Control Information Item (ACI)**.

This appendix contains these topics:

# Schema for orclACI

The access control directive defined by the user attribute orclACI has the following schema:

```
OrclACI:
{ object_identifier NAME 'orclACI' DESC 'Stores an inheritable ACI' EQUALITY
accessDirectiveMatch SYNTAX 'accessDirectiveDescription'  USAGE
'directoryOperation'}
```

accessDirectiveDescription has the following BNF:

```
<accessDirectiveDescription>
                  ::= access to <object> [by <subject> ( <accessList> )]+

<object> ::= [attr <EQ-OR-NEQ> (<attrList>) | entry] [filter=(<ldapFilter>)]

<subject> ::= <entity> [<BindMode>]

<entity> ::= * | self | dn="<regex>" | dnAttr=(<dn_attribute>) | group="<dn>"

<BindMode> ::= | BindMode = Simple
               | BindMode = SSLNoauth
               | BindMode = SSLOneway
               | BindMode = SSLTwoway

<accessList> ::= <access> | <access>, <accessList>

<access> ::= none | compare | search | browse | read | selfwrite | write | add
|delete | nocompare | nosearch | nobrowse |noread | noselfwrite | nowrite |
noadd | nodelete

<attrList> ::=  *| <attribute name> | <attribute name>,<attrList>

<EQ-OR-NEQ> ::=  = | !=

<regex> ::= <dn> | *,<dn_of_any_subtree_root>
```

---

**Note:** The regular expression defined above is not meant to match any arbitrary expression. The syntax only allows expressions where the wild card is followed by a comma and a valid DN. The latter DN denoted by *<dn_of_any_subtree_root>* is intended to specify the root of some subtree.

---

# Schema for orclEntryLevelACI

The entry level access control directive defined by the user attribute orclEntryLevelACI has the following schema:

```
"orclEntryLevelACI":
{ object_identifier NAME 'orclEntryLevelACI' DESC 'Stores entry level ACL
Directive'
EQUALITY accessDirectiveMatch SYNTAX 'orclEntryLevelACIDescription'
USAGE 'directoryOperation' }


<orclEntryLevelACIDescription>
::= access to <object> [by <subject> ( <accessList> )]+
```

# E

# Schema Elements

This appendix briefly lists different schema elements supported in the Oracle Internet Directory. Most of these elements are used as defined by the ldapext and ASID working groups of the Internet Engineering Task Force (IETF).

**See Also:** The following URLs on the World Wide Web:

- `http://www.ietf.org` for the IETF home page
- `http://www.ietf.org/html.charters/ldapext-charter.html` for the ldapext charter and LDAP drafts)
- `http://ietf.org/html.charters/asid-charter.html` for the ASID charter and LDAP drafts
- `http://www.ietf.org/html.charters/ldup-charter.html` for the LDUP charter and drafts
- `http://www.iana.org`, the Internet Assigned Numbers Authority home page, for information about object identifiers

This appendix contains these topics:

# IETF Requests for Comments (RFCs) Enforced by Oracle Internet Directory

Oracle Internet Directory enforces the following Requests for Comments (RFCs) of the Internet Engineering Task Force (IETF):

| RFC | Title | URL |
|---|---|---|
| 1777 | Lightweight Directory Access Protocol | http://www.ietf.org/rfc/rfc1777.txt |
| 1778 | The String Representation of Standard Attribute Syntaxes | http://www.ietf.org/rfc/rfc1778.txt |
| 1779 | A String Representation of Distinguished Names | http://www.ietf.org/rfc/rfc1779.txt |
| 1960 | A String Representation of LDAP Search Filters | http://www.ietf.org/rfc/rfc1960 |
| 2079 | Definition of an X.500 Attribute Type and an Object Class to Hold Uniform Resource Identifiers (URIs) | http://www.ietf.org/rfc/rfc2079.txt |
| 2247 | Using Domains in LDAP/X.500 Distinguished Names | http://www.ietf.org/rfc/rfc2247.txt |
| 2251 | Lightweight Directory Access Protocol (v3) | http://www.ietf.org/rfc/rfc2251.txt |
| 2252 | Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions | http://www.ietf.org/rfc/rfc2252.txt |
| 2253 | Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names | http://www.ietf.org/rfc/rfc2253.txt |
| 2254 | The String Representation of LDAP Search Filters | http://www.ietf.org/rfc/rfc2254.txt |
| 2255 | The LDAP URL Format | http://www.ietf.org/rfc/rfc2255.txt |
| 2256 | A Summary of the X.500(96) User Schema for use with LDAPv3 | http://www.ietf.org/rfc/rfc2256.txt |

# IETF Drafts Enforced by Oracle Internet Directory

Oracle Internet Directory enforces the following two drafts of the IETF:

Draft: "Definition of the inetOrgPerson LDAP Object Class"

URL: http://ietf.org/internet-drafts/draft-smith-ldap-inetorgperson-03.txt

Draft: "Referrals and Knowledge References in LDAP Directories"

URL: http://www.ietf.org/internet-drafts/draft-ietf-ldapext-knowledge-00.txt

# Proprietary Oracle Internet Directory Schema Elements

Oracle Internet Directory's proprietary schema includes attributes and object classes in these categories:

- Access Control
- Replication
- Oracle Internet Directory Configuration
- SSL
- Audit Log
- Configuration Set Entry Attributes

In addition, Oracle Internet Directory installation includes schema elements that enable specific Oracle products to use Oracle Internet Directory. For information about these schema elements, see the documentation for the specific Oracle product.

## Access Control

Attributes `orclEntryLevelACI, orclACI`

Object Class `orclPrivilegeGroup`

### Replication

Attributes: `orclGUID, changeNumber changeType, changes, orclParentGUID, server, supplier, consumer, orclReplBindDN, orclReplBindPassword, changeLog, changeStatus, orclChangeRetryCount, orclPurgeSchedule, orclDirReplGroupAgreement, orclAgreementId, orclSupplierReference,orclConsumerReference, orclReplicationProtocol, orclUpdateSchedule, targetDN, orclExcludedNamingcontexts, orclDirReplGroupDSAs`

Object class: changeLogEntry, changeStatusEntry, orclReplAgreementEntry

### Oracle Internet Directory Configuration

Attributes: `orclDebugLevel, orclMaxCC, orclDBType, orclSuffix, orclDITRoot, orclSuName, orclSuPassword, orclSizeLimit, orclTimeLimit, orclGuName, orclGuPassword, orclServerProcs, orclconfigsetnumber, orclhostname, orclIndexedAttribute, orclCatalogEntryDN, orclServerMode, orclPrName, orclPrPassword, orclUseEncrypt, orclDirectoryVersion`

Object class: `subconfig, orclConfigSet, orclLDAPSubConfig, orclREPLSubConfig, orclcontainerOC, subregistry, orclLDAPInstance, orclREPLInstance, orclIndexOC, orcleventLog, orclEvents`

### SSL

**Note:** These attribute values are stored as part of configuration entries.

Attributes: `orclsslAuthentication, orclsslEnable, 'orclsslWalletURL, orclsslWalletPasswd, orclsslPort, orclsslVersion`

### Audit Log

| | |
|---|---|
| Attributes | `orclServerEvent, orcleventtype, orclauditattribute, orclauditmessage, orcleventtime, orcluserdn, orclSequence, orclAuditLevel, orclOpResult` |
| Object class | `OrclAuditOC` |

### Configuration Set Entry Attributes

The following table lists and describes the entire set of configuration set entry attributes that are used to configure an instance of a directory server.

| Parameter | Description |
|---|---|
| `orcldebuglevel` | Debug level associated with this instance of the server. The default for configset0 is 0. The range is 0 to 65535. |
| `orclmaxcc` | Maximum number of concurrent database connections. The default for configset0 is 10. You cannot use a negative value for this attribute. |
| `orclserverprocs` | Number of server processes to start. The default for configset0 is 1. You cannot use a negative value for this attribute. |
| `orclsslport` | SSL mode default port (default 636). When you run the directory in the secure mode, it listens at default port 636 and accepts only SSL-based TCP/IP connections. (When you run the directory in the normal mode, it listens at default port 389, accepting normal TCP/IP connections.) You might want to change this port when you add multiple LDAP server instances. |
| `orclnonsslport` | Non-SSL mode default port (default 389). |
| `orclsslenable` | Flag for toggling SSL on and off. You would want to toggle this flag when you use different instances of the same server for either SSL or non-SSL. You may use either of the following two values:<br>■ 0 = disables SSL (default in configuration set0)<br>■ 1 = enables SSL<br>The default is 0. |

| Parameter | Description |
|---|---|
| orclsslauthentication | Flag, with values of 1, 32, or 64, for specifying the type of authentication you elect to use for each instance of the Oracle directory server. The default value, 1, specifies no authentication. You can run different values concurrently for different instances. Values of one-way and two-way authentication require wallets. You may use one of the following three values:<br>■ 1 = no SSL authentication<br>■ 32 = one-way SSL authentication (the server sends its certificate to the client)<br>■ 64 = two-way SSL authentication (client and server send certificates to each other) |
| orclsslwalleturl | Sets the location of the Oracle wallet. You initially set this value when you create the wallet. If you elect to change the location of the Oracle wallet, you must change this parameter. You must set the wallet location on both the client and the server. For example, on Solaris, you could set this parameter as follows:<br>`orclsslwalleturl=file:/Home/my_dir/`<br>On Windows NT, you could set this parameter as follows:<br>`file:Home\my_dir\` |
| orclsslwalletpasswd | Password used by the server to open its wallet. You initially set this value when you create the wallet. If you elect to change the wallet password, you must change this parameter. You must set the wallet password on both the client and the server. |
| orclsslversion | SSL version. The default is 3. |

**See Also:**

- "Setting Debug Logging Levels by Using the OID Control Utility" on page 5-23 for information on debug levels
- Appendix C, "Using Oracle Wallet Manager" for information on setting the location of the Oracle Wallet and the Oracle Wallet password

# LDAP Syntax

Syntax defines the type of values that an attribute can hold. Oracle Internet Directory recognizes most of the syntax specified in RFC 2252, that is, it allows you to associate most of the syntax described in that document with an attribute. In addition to recognizing most LDAP syntax, Oracle Internet Directory enforces some LDAP syntax.

This section covers topics in the following subsections:



## LDAP Syntax Enforced by Oracle Internet Directory

Oracle Internet Directory enforces LDAP syntax for the following:

- DN
- Facsimile Telephone Number
- OID (object identifier)
- Telephone Number

---

**Note:** The values you specify for these attributes must conform to the syntax specified in RFC 2252.

---

## Commonly Used LDAP Syntax Recognized by Oracle Internet Directory

The following LDAP syntax is more commonly used:

| | |
|---|---|
| Attribute Type Description | Numeric String |
| Boolean | Object Class Description |
| Certificate | Octet String |
| Directory String | OID |
| DN | Presentation Address |
| Facsimile Telephone Number | Printable String |
| INTEGER | Telephone Number |
| JPEG | UTC Time |
| Name And Optional UID | |

## Additional LDAP Syntax Recognized by Oracle Internet Directory

In addition to the commonly used LDAP syntax defined above, Oracle Internet Directory recognizes LDAP syntax for the following:

Access Point
ACI Item
Audio
Binary
Bit String
Certificate List
Certificate Pair
Country String
Data Quality Syntax
Delivery Method
DIT Content Rule Description
DIT Structure Rule Description
DL Submit Permission
DSA Quality Syntax
DSE Type
Enhanced Guide
Fax
Generalized Time
Guide
IA5 String
LDAP Schema Definition
LDAP Schema Description
LDAP Syntax Description
Mail Preference
Master And Shadow Access Points
Matching Rule
Matching Rule Use Description
MHS OR Address
Modify Rights
Name Form Description
Object Class Description
Octet String
Other Mailbox
Postal Address
Protocol Information
Substring Assertion
Subtree Specification
Supplier And Consumer
Supplier Information
Supplier Or Consumer
Supported Algorithm
Teletex TerminalIdentifier
Telex Number

### Size of Attribute Values

Syntax does not put any specific size constraint on attribute values. You can, however, use syntax to specify the size of the attribute value. Oracle Internet Directory does not enforce the 'len' characteristics on the attribute.

For example, to limit an attribute foo to a size of 64, you would define the attribute as follows:

(*object_identifier_of_attribute* NAME 'foo' EQUALITY caseIgnoreMatch SYNTAX '*object_identifier_of_syntax*{64}')

**See Also:** Section 4.1.6 f of RFC2251 for more information on Attribute Value. You can find this RFC at the following URL: http://www.ietf.org/rfc/rfc2251.txt.

## Matching Rules

Oracle Internet Directory recognizes the following matching rules definitions in the schema.

accessDirectiveMatch
bitStringMatch
caseExactMatch
caseExactIA5Match
caseIgnoreIA5Match
caseIgnoreListMatch
caseIgnoreMatch
caseIgnoreOrderingMatch
distinguishedNameMatch
generalizedTimeMatch
generalizedTimeOrderingMatch
IntegerMatch
numericStringMatch
objectIdentifierFirstComponentMatch
ObjectIdentifierMatch
OctetStringMatch
presentationAddressMatch
protocolInformationMatch
telephoneNumberMatch
uniqueMemberMatch

Of the matching rules in the previous list, Oracle Internet Directory actually enforces the following when it compares attribute values:

DistinguishedNameMatch

caseExactMatch

caseIgnoreMatch

numericStringMatch

IntegerMatch

telephoneNumberMatch

# F

# Migrating Data from Other LDAP-Compliant Directories

This appendix describes the steps to migrate data from LDAP v3-compatible directories into Oracle Internet Directory.

This appendix contains these topics:

## About the Data Migration Process

This method uses the established LDIF file format for LDAP v3 flat file representation of application data and metadata. LDIF is the IETF-sanctioned ASCII interchange format for representing LDAP v3 directory data as a file. All LDAP v3-compatible servers should be able to export their contents into one or more LDIF files representing the directory information tree at the time of export. However, not all LDIF files are created equal: Certain proprietary attributes or metadata may or may not be included in a give product's LDIF output. As a result, there are some additional steps required before importing an LDIF file back into Oracle Internet Directory when using bulkload or ldapadd.

> **See Also:** `//ftp.isi.edu/in-notes/rfc2849.txt` for more on the LDIF Technical Specification

## Migrating Data

This section contains these topics:

- Task 1: Export Data from the Non-Oracle Internet Directory Server into LDIF File Format
- Task 2: Analyze the LDIF User Data for Any Required Schema Additions Referenced in the LDIF Data
- Task 3: Extend the Schema in Oracle Internet Directory
- Task 4: Remove Any Proprietary Directory Data from the LDIF File
- Task 5: Remove Operational Attributes from the LDIF File
- Task 6: Remove Incompatible userPassword Attribute Values from the LDIF File
- Task 7: Run the bulkload.sh -check Mode and Determine Any Remaining Schema Violations or Duplication Errors

## Task 1: Export Data from the Non-Oracle Internet Directory Server into LDIF File Format

See the vendor-supplied documentation for instructions. If flags or options exist for exporting data from the foreign directory, be sure to select the method that:

- Produces LDIF output with the least amount of proprietary information included
- Provides maximum conformance to the IETF Request for Comments 2849 mentioned on page F-2

## Task 2: Analyze the LDIF User Data for Any Required Schema Additions Referenced in the LDIF Data

Any attributes not found in the Oracle Internet Directory base schema require extension of the Oracle Internet Directory base schema prior to the importation of the LDIF file. Some directories may support the use of configuration ("conf") files for defining extensions to their base schema (Oracle Internet Directory does not). If you have a configuration file you can use it as a guideline for extending the base schema in Oracle Internet Directory in "Task 3: Extend the Schema in Oracle Internet Directory".

## Task 3: Extend the Schema in Oracle Internet Directory

See the chapter on managing the directory schema in *Oracle Internet Directory Administrator's Guide* for tips on how to extend directory schema in Oracle Internet Directory. You can do this by using either Oracle Directory Manager or command-line tools.

## Task 4: Remove Any Proprietary Directory Data from the LDIF File

Certain elements of the LDAP v3 standard have not yet been formalized, such as Access Control Information (ACI) attributes. As a result, various directory vendors implement ACI policy objects in ways that do not "port" across vendor installations.

After the basic entry data has been imported from the "sanitized" LDIF file, you must explicitly re-apply security policies in the Oracle Internet Directory environment. You can do this by using either Oracle Directory Manager, or command line tools and LDIF files containing the desired Access Control Policy information.

There may be other proprietary metadata, representing areas outside the area of access control, that you should remove as well. A thorough understanding of the various IETF RFCs can help you determine which directory metadata is proprietary to a given vendor and which is standards-compliant, and thus portable by way of an LDIF file.

## Task 5: Remove Operational Attributes from the LDIF File

Two of the standard LDAP v3 operational attributes, namely, `creatorsName`, `createTimestamp`, `modifiersName`, and `modifyTimestamp` are automatically generated by Oracle Internet Directory whenever entries are created or imported. It is not possible to instantiate these values from existing directory data, for example by using LDIF file importation. Therefore you should remove these attributes from the file before attempting to import.

## Task 6: Remove Incompatible userPassword Attribute Values from the LDIF File

Oracle Internet Directory release 2.1.1 supports the following `userPassword` attribute hash algorithms:

- **MD4**
- **MD5**
- No encryption
- **SHA**
- **UNIX Crypt**

The `userPassword` attribute hash values used by some vendor products are not compatible with Oracle Internet Directory. As a result, all lines corresponding to the `userPassword` attribute and value should be pruned from the LDIF data file unless they are represented in plain text or contain no value. After importation of the LDIF data, you must re-enter manually or upload hashed `userPassword` information separately into the directory.

## Task 7: Run the bulkload.sh -check Mode and Determine Any Remaining Schema Violations or Duplication Errors

Before generating and loading an LDIF file, always perform check mode on it by using the bulkload utility. The bulkload output reports any inconsistencies in the data.

# G

# Troubleshooting

This appendix explains typical problems that you could encounter while running or installing the Oracle Internet Directory. It contains these topics:

- Installation Errors
- Administration Error Messages and Causes

# Installation Errors

During installation and configuration of the Oracle8*i* database server, you must select the character set UTF-8. If you select any other character set, the directory server will not function properly.

# Administration Error Messages and Causes

This section contains a list of all the Oracle directory server error messages that you can encounter. Each message is followed by its most probable causes.

This section contains these topics:



## Oracle Database Server Error Due to Schema Modifications

**ORA-1562**

**Cause:** If you attempt to add more schema components than can fit in the rollback segment space, you will encounter this error and the modifications will not commit. To solve this, increase the size of the rollback segments in the database server.

## Standard Error Messages Returned from Oracle Directory Server

The following are standard error messages. Oracle Internet Directory also returns other messages listed and described in "Additional Error Messages" on page G-6.

**00—LDAP_SUCCESS**

**Cause:** The operation was successful.

**01—LDAP_OPERATIONS_ERROR**

**Cause:** General errors encountered by the server when processing the request.

**02—LDAP_PROTOCOL_ERROR**

**Cause:** The client request did not meet the LDAP protocol requirements, such as format or syntax. This can occur in the following situations:

- Server encounters a decoding error while parsing the incoming request
- The request is an add or modify request that specifies the addition of an attribute type to an entry but no values specified

- Error reading SSL credentials
- An unknown type of modify operation is specified (other than LDAP_MOD_ADD, LDAP_MOD_DELETE, and LDAP_MOD_REPLACE)
- Unknown search scope

**03—LDAP_TIMELIMIT_EXCEEDED**

**Cause:** Search took longer than the time limit specified. If you have not specified a time limit for the search, Oracle Internet Directory uses a default time limit of one hour.

**04—LDAP_SIZELIMIT_EXCEEDED**

**Cause:** More entries match the search query than the size limit specified. If you have not specified a size limit for the search, Oracle Internet Directory uses a default size limit.

**05—LDAP_COMPARE_FALSE**

**Cause:** Presented value is not the same as the one in the entry.

**06—LDAP_COMPARE_TRUE**

**Cause:** Presented value is same as the one in the entry.

**07—LDAP_STRONG_AUTH_NOT_SUPPORTED**

**Cause:** Bind method is not supported by the server.

**08—LDAP_STRONG_AUTH_REQUIRED**

**Cause:** Strong authentication is required. Oracle Internet Directory does not return this message at the present time.

**09—LDAP_PARTIAL_RESULTS**

**Cause:** Server returned a referral.

**10—LDAP_REFERRAL**

**Cause:** Server returned a referral.

**11—LDAP_ADMINLIMIT_EXCEEDED**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**12—LDAP_UNAVAILABLE_CRITICALEXTENSION**

**Cause:** Specified request is not supported

**16—LDAP_NO_SUCH_ATTRIBUTE**

**Cause:** Attribute does not exist in the entry specified in the request.

**17—LDAP_UNDEFINED_TYPE**

**Cause:** Specified attribute type is undefined in the schema.

**18—LDAP_INAPPROPRIATE_MATCHING**

**Cause:** Specified matching rule is inappropriate for the attribute type. Oracle Internet Directory does not return this message at the present time.

**19—LDAP_CONSTRAINT_VIOLATION**

**Cause:** The value in the request violated certain constraints.

**20—LDAP_TYPE_OR_VALUE_EXISTS**

**Cause:** Duplicate values specified for the attribute.

**21—LDAP_INVALID_SYNTAX**

**Cause:** Specified *attribute* syntax is invalid. In a search, the *filter* syntax is invalid.

**32—LDAP_NO_SUCH_OBJECT**

**Cause:** The base specified for the operation does not exist.

**33—LDAP_ALIAS_PROBLEM**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**34—LDAP_INVALID_DN_SYNTAX**

**Cause:** Error in the DN syntax.

**35—LDAP_IS_LEAF**

**Cause:** The entry is a leaf (terminal entry). Oracle Internet Directory does not return this message at the present time.

**36—LDAP_ALIAS_DEREF_PROBLEM**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**48—LDAP_INAPPROPRIATE_AUTH**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**49—LDAP_INVALID_CREDENTIALS**

**Cause:** Bind failed because the credentials are not correct.

**50—LDAP_INSUFFICIENT_ACCESS**

**Cause:** The client does not have access to perform this operation.

**51—LDAP_BUSY**

**Cause:** Server cannot accept any more client connections. Oracle Internet Directory does not return this message at the present time.

**52—LDAP_UNAVAILABLE**

**Cause:** Cannot contact the server at all. Oracle Internet Directory does not return this message at the present time.

**53—LDAP_UNWILLING_TO_PERFORM**

**Cause:** General error, or server is in read-only mode.

**54—LDAP_LOOP_DETECT**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**64—LDAP_NAMING_VIOLATION**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**65—LDAP_OBJECT_CLASS_VIOLATION**

**Cause:** A change to the entry violates the objectclass definition.

**66— LDAP_NOT_ALLOWED_ON_NONLEAF**

**Cause:** The entry to be deleted has children.

**67—LDAP_NOT_ALLOWED_ON_RDN**

**Cause:** Cannot perform the operation on RDN attributes—for example, you cannot delete the RDN attribute of the entry.

**68—LDAP_ALREADY_EXISTS**

**Cause:** Duplicate ADD condition.

**69—LDAP_NO_OBJECT_CLASS_MODS**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**70—LDAP_RESULTS_TOO_LARGE**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**80—LDAP_OTHER**

**Cause:** Oracle Internet Directory does not return this message at the present time.

**81—LDAP_SERVER_DOWN**

**Cause:** Can't contact LDAP server. This message is returned from the SDK.

**82—LDAP_LOCAL_ERROR**

**Cause:** The client encountered an internal error. This message is returned from the client SDK.

**83—LDAP_ENCODING_ERROR**

**Cause:** The client encountered an error in encoding the request. This message is returned from the SDK.

**84—LDAP_DECODING_ERROR**

**Cause:** The client encountered an error in decoding the request. This message is returned from the SDK.

**85—LDAP_TIMEOUT**

**Cause:** Client encountered the time-out specified for the operation. This message is returned from the SDK.

**86—LDAP_AUTH_UNKNOWN**

**Cause:** Authentication method is unknown to the client SDK.

**87—LDAP_FILTER_ERROR**

**Cause:** Bad search filter

**88—LDAP_USER_CANCELLED**

**Cause:** User cancelled operation

**89—LDAP_PARAM_ERROR**

**Cause:** Bad parameter to an LDAP routine

**90—LDAP_NO_MEMORY**

**Cause:** Out of memory

## Additional Error Messages

These messages do not display error codes.

The Oracle Internet Directory application replaces the *parameter* tag seen in some of the messages below with the appropriate run-time value.

**%s attribute not found.**

**Cause:** The particular attribute type is not defined in the schema.

**<parameter> not found for attribute <parameter>.**

**Cause:** Value not found in the attribute. (ldapmodify)

**Admin domain does not contain schema information for objectclass <parameter>.**

**Cause:** The object class specified in the request is not present in the schema.

**Attempted to add a Class with oid <parameter> taken by other class.**

**Cause:** Duplicate object identifier specified. (schema modification)

**Attribute <parameter> already in use.**

**Cause:** Duplicate attribute name. (schema modification)

**Attribute <parameter> has syntax error.**

**Cause:** Syntax error in the attribute name definition. (schema modification)

**Attribute <parameter> is not supported in the schema.**

**Cause:** Attribute not defined. (all operations)

**Attribute <parameter> is single valued.**

**Cause:** Attribute is single-valued. (ldapadd & ldapmodify)

**Attribute <parameter> not present in the entry.**

**Cause:** This attribute does not exist in the entry. (ldapmodify)

**Bad attribute definition.**

**Cause:** Syntax error in attribute definition. (schema modification)

**Currently Not Supported**

**Cause:** The version of LDAP request is not supported by this server.

**Entry to be deleted not found.**

**Cause:** DN specified in the delete operation not found.

**Entry to be modified not found**

**Cause:** The entry specified in the request is not found.

**Error encountered while adding <parameter> to the entry**

**Cause:** Returned when modify add operation is invoked. A possible cause is that the system resource is unavailable.

**Error encountered while encrypting an attribute value.**

**Cause:** Error in encrypting user password. (all operations)

**Error in DN Normalization.**

**Cause:** DN specified is invalid. Syntax error encountered in parsing the DN. (all operations)

**Error in hashing <parameter> attribute.**

**Cause:** Error in creating hash entry for the attribute. (schema modification)

**Error in hashing <parameter> objectclass.**

**Cause:** Error in creating hash entry for the objectclass. (schema modification)

**Error in Schema hash creation.**

**Cause:** Error while creating hash table for schema. (schema modification)

**Error replacing <parameter>.**

**Cause:** Error in replacing this attribute. (ldapmodify)

**Error while normalizing value for attribute <parameter>.**

**Cause:** Error in normalizing value for the attribute. (all operations)

**Failed to find <parameter> in mandatory or optional attribute list.**

**Cause:** Attribute specified does not exist in either the mandatory or optional attribute list as required by the object class(es).

**Function Not Implemented**

**Cause:** The feature / request is currently not supported.

**INVALID ACI is <parameter>**

**Cause:** The particular ACI you specified in a request is invalid.

**Mandatory attribute <parameter> is not defined in Admin Domain <parameter>.**

**Cause:** MUST refers to attribute not defined. (schema modification)

**Mandatory Attribute missing.**

**Cause:** The mandatory attribute for the particular entry is missing, as required by the particular object class.

**Matching rule, <parameter>, not defined.**

**Cause:** Matching rule not defined in the server. (schema modification)

**MaxConn Reached**

**Cause:** The maximum number of concurrent connections to the LDAP server has been reached.

**Modifying the Naming attribute for the entry without modifying the DN.**

**Cause:** Cannot modify the naming attributes using ldap_modify. A naming attribute, such as *cn* is an element in the DN.

**New Parent not found.**

**Cause:** New parent specified in modifydn operation does not exist.(ldapmodifydn)

**Object already exists.**

**Cause:** Duplicate entry. (ldapadd and ldapmodifydn)

**Object ID <parameter> already in use.**

**Cause:** Duplicate object identifier specified. (schema modification)

**Objectclass <parameter> already in use. m**

**Cause:** Duplicate Objectclass name. (schema modification)

**Objectclass attribute missing.**

**Cause:** The objectclass attribute is missing for this particular entry.

**OID <parameter> has syntax error.**

**Cause:** syntax error in the object identifier definition. (schema modification)

**One of the attributes in the entry has duplicate value**

**Cause:** You entered two values for the same attribute in the entry you are creating.

**Operation not allowed on the <parameter>.**

**Cause:** Operation not allowed on this entry. (modify, add, and delete)

**Operation not allowed on the DSE Entry.**

**Cause:** Can't do this operation on DSE entry. (delete)

**Optional attribute <parameter> is not defined in Admin Domain <parameter>.**

**Cause:** MAY refers to attribute not defined. (schema modification)

**Parent entry not found in the directory.**

**Cause:** Parent entry does not exist. (ldapadd and perhaps ldapmodifydn)

**Super object <parameter> is not defined in Admin Domain <parameter>.**

**Cause:** SUP types refer to non-existing class. (schema modification)

**Super type undefined.**

**Cause:** SUP type does not exist. (schema modification)

**Super user addition not permitted.**

**Cause:** Cannot create super user entry. (ldapadd)

**Syntax, <parameter>, not defined.**

**Cause:** Syntax not defined in the server. (schema modification)

**The attribute or the value specified in the RDN does not exist in the entry.**

**Cause:** AVA specified as the RDN does not exist in the entry. (ldapadd)

**Unknown search scope**

**Cause:** The search scope specified in the LDAP request is not recognized.

**Version Not Supported**

**Cause:** The version of the LDAP request is not supported by this server.

# Glossary

### Access Control Information Item (ACI)

An attribute that determines who has what type of access to what directory data. It contains a set of rules for structural access items, which pertain to entries, and content access items, which pertain to attributes. Access to both structural and content access items may be granted to one or more users or groups.

### Access Control List (ACL)

The group of access directives that you define. The directives grant levels of access to specific data for specific clients, or groups of clients, or both.

### Access Control Policy Point

An entry that contains security directives that apply downward to all entries at lower positions in the **directory information tree (DIT)**.

### ACI

See **Access Control Information Item (ACI)**.

### ACL

See **Access Control List (ACL)**.

### ACP

See **Access Control Policy Point**.

### Advanced Symmetric Replication (ASR)

A feature in Oracle8*i* that allows database tables to be kept synchronized across two Oracle databases.

### API

See **Application Program Interface**.

### Application Program Interface

Programs to access the services of a specified application. For example, LDAP-enabled clients access directory information through programmatic calls available in the LDAP API.

### administrative area

A subtree on a directory server whose entries are under the control (schema, ACL, and collective attributes) of a single administrative authority.

### anonymous authentication

The process by which the directory authenticates a user without requiring a user name and password combination. Each anonymous user then exercises the privileges specified for anonymous users.

### API

See **Application Program Interface**.

### ASR

See **Advanced Symmetric Replication (ASR)**.

### attribute

An item of information that describes some aspect of an entry. An entry comprises a set of attributes, each of which belongs to an **object class**. Moreover, each attribute has both a *type*, which describes the kind of information in the attribute, and a *value*, which contains the actual data.

### authentication

The process of verifying the identity of a user, device, or other entity in a computer system, often as a prerequisite to allowing access to resources in a system.

### authorization

Permission given to a user, program, or process to access an object or set of objects.

### binding

The process of authenticating to a directory.

### certificate

An ITU x.509 v3 standard data structure that securely binds an identity to a public key. A certificate is created when an entity's public key is signed by a trusted identity: a **certificate authority (CA)**. This certificate ensures that the entity's information is correct and that the public key actually belongs to that entity.

### certificate authority (CA)

A trusted third party that certifies that other entities—users, databases, administrators, clients, servers—are who they say they are. The certificate authority verifies the user's identity and grants a certificate, signing it with the certificate authority's private key.

### certificate chain

An ordered list of certificates containing an end-user or subscriber certificate and its certificate authority certificates.

### change logs

A database that records changes made to a directory server.

### cipher suite

In SSL, a set of authentication, encryption, and data integrity algorithms used for exchanging messages between network nodes. During an SSL handshake, the two nodes negotiate to see which cipher suite they will use when transmitting messages back and forth.

### cold backup

The procedure to add a new **DSA** to an existing replicating system by using the database copy procedure.

### concurrent clients

The total number of clients that have established a session with Oracle Internet Directory.

### concurrent operations

The number of operations that are being executed on the directory from all of the concurrent clients. Note that this is not necessarily the same as the concurrent clients, because some of the clients may be keeping their sessions idle.

### configset

See **configuration set entry**.

### configuration set entry

A directory entry holding the configuration parameters for a specific instance of the directory server. Multiple configuration set entries can be stored and referenced at run-time. The configuration set entries are maintained in the subtree specified by the subConfigsubEntry attribute of the DSE, which itself resides in the associated **directory information base (DIB)** against which the servers are started.

### connect descriptor

A specially formatted description of the destination for a network connection. A connect descriptor contains destination service and network route information.

The destination service is indicated by using its service name for Oracle8*i* release 8.1 database or its Oracle System Identifier (SID) for Oracle release 8.0 or version 7 databases. The network route provides, at a minimum, the location of the listener through use of a network address.

### consumer

A directory server that is the destination of replication updates. Sometimes called a slave.

### context prefix

The **DN** of the root of a **directory naming context**.

### data integrity

The guarantee that the contents of the message received were not altered from the contents of the original message sent.

### decryption

The process of converting the contents of an encrypted message (ciphertext) back into its original readable format (plaintext).

### default knowledge reference

A **knowledge reference** that is returned when the base object is not in the directory, and the operation is performed in a naming context not held locally by the server. A default knowledge reference typically sends the user to a server that has more knowledge about the directory partitioning arrangement.

### DES

Data Encryption Standard, a block cipher developed by IBM and the U.S. government in the 1970's as an official standard.

### DIB

See **directory information base (DIB)**.

### directory information base (DIB)

The complete set of all information held in the directory. The DIB consists of entries that are related to each other hierarchically in a **directory information tree (DIT).**

### directory information tree (DIT)

A hierarchical tree-like structure consisting of the DNs of the entries.

### directory naming context

See **naming context**.

### directory replication group (DRG)

The directory servers participating in a replication agreement.

### directory system agent (DSA)

The X.500 term for a directory server.

### distinguished name (DN)

The unique name of a directory entry. It comprises all of the individual names of the parent entries back to the root.

### DIT

See **directory information tree (DIT)**.

### DN

See **distinguished name (DN)**.

### DRG

See **directory replication group (DRG)**.

### DSA

See **directory system agent (DSA)**.

### DSE

DSA specific entries. Different DSAs may hold the same DIT name, but have different contents. That is, the contents can be specific to the DSA holding it. A DSE is an entry with contents specific to the DSA holding it.

### encryption

The process of disguising the contents of a message and rendering it unreadable (ciphertext) to anyone but the intended recipient.

### entry

The building block of a directory, it contains information about an object of interest to directory users.

### failover

The process of failure recognition and recovery.

### filter

A method of qualifying data, usually data that you are seeking. Filters are always expressed as DNs, for example: `cn=susie smith, o=acme, c=us`.

### global unique identifier (GUID)

In a multi-master replication environment, an entry replicated on multiple nodes has the same DN on each node. However, even though it has the same DN, it is assigned a different GUID on each node. For example, the same DN can be replicated on both node1 and node2, but the GUID for that DN as it resides on node1 would be different from the GUID for that DN on node2.

### guest user

One who is not an anonymous user, and, at the same time, does not have a specific user entry.

### GUID

See **global unique identifier (GUID)**.

### handshake

A protocol two computers use to initiate a communication session.

### inherit

When an object class has been derived from another class, it also derives, or inherits, many of the characteristics of that other class. Similarly, an attribute subtype inherits the characteristics of its supertype.

### instance

See **server instance**.

### integrity

The guarantee that the contents of the message received were not altered from the contents of the original message sent.

### Internet Message Access Protocol (IMAP)

A protocol allowing a client to access and manipulate electronic mail messages on a server. It permits manipulation of remote message folders, also called mailboxes, in a way that is functionally equivalent to local mailboxes.

### key

A string of bits used widely in cryptography, allowing people to encrypt and decrypt data; a key can be used to perform other mathematical operations as well. Given a cipher, a key determines the mapping of the plaintext to the ciphertext.

### knowledge reference

The access information (name and address) for a remote **DSA** and the name of the **DIT** subtree that the remote DSA holds. Knowledge references are also called referrals.

### latency

The time a client has to wait for a given directory operation to complete.

### LDAP

See **Lightweight Directory Access Protocol (LDAP)**.

### Lightweight Directory Access Protocol (LDAP)

A standard, extensible directory access protocol. It is a common language that LDAP clients and servers use to communicate. The framework of design conventions supporting industry-standard directory products, such as the Oracle Internet Directory.

### LDAP Data Interchange Format (LDIF)

The set of standards for formatting an input file for any of the LDAP command line utilities.

### master definition site (MDS)

In replication, a Master Definition Site is the Oracle Internet Directory database from which the administrator runs the configuration scripts.

### master site

In replication, a master site is any site other than the Master Definition Site that participates in LDAP replication.

### matching rule

In a search or compare operation, determines equality between the attribute value sought and the attribute value stored. For example, matching rules associated with the `telephoneNumber` attribute could cause "(650) 123-4567" to be matched with either "(650) 123-4567" or "6501234567" or both. When you create an attribute, you associate a matching rule with it.

### MD4

A one-way hash function that produces a 128-bit hash, or message digest. If as little as a single bit value in the file is modified, the MD4 checksum for the file will change. Forgery of a file in a way that will cause MD4 to generate the same result as that for the original file is considered extremely difficult.

### MD5

An improved version of MD4.

### MDS

See **master definition site (MDS)**.

### MTS

See **multi-threaded server (MTS)**.

### multi-threaded server (MTS)

A server that is configured to allow many user processes to share very few server processes, so the number of users that can be supported is increased. With MTS configuration, many user processes connect to a dispatcher. The dispatcher directs multiple incoming network session requests to a common queue. An idle shared server process from a shared pool of server processes picks up a request from the queue. This means a small pool of server processes can server a large amount of clients. Contrast with dedicated server.

### naming attribute

A specialized attribute that holds values for different types of **RDN**. A naming attribute is identifiable by its mnemonic label, usually `cn`, `sn`, `ou`, `o`, `c`, and so on. For example, the naming attribute `c` is the mnemonic for the naming attribute `country`, and it holds the RDN for specific country values.

### naming context

A subtree that resides entirely on one server. It must be contiguous, that is, it must begin at an entry that serves as the top of the subtree, and extend downward to either leaf entries or **knowledge reference**s (also called referrals) to subordinate naming contexts. It can range in size from a single entry to the entire DIT.

### Net8

The foundation of the Oracle family of networking products, allowing services and their client applications to reside on different computers and communicate. The main function of Net8 is to establish network sessions and transfer data between a client application and a server. Net8 is located on each computer in the network. Once a network session is established, Net8 acts as a data courier for the client and the server.

### net service name

A simple name for a service that resolves to a connect descriptor. Users initiate a connect request by passing a user name and password along with a net service name in a connect string for the service to which they wish to connect:

```
CONNECT username/password@net_service_name
```

Depending on your needs, net service names can be stored in a variety of places, including:

- Local configuration file, tnsnames.ora, on each client
- Directory server
- Oracle Names server
- External naming service, such as NDS, NIS or CDS

### object class

A named group of attributes. When you want to assign attributes to an entry, you do so by assigning to that entry the object classes that hold those attributes.

All objects associated with the same object class share the same attributes.

### OID Control Utility

A command-line tool for issuing run-server and stop-server commands. The commands are interpreted and executed by the **OID Monitor** process.

### OID Database Password Utility

The utility used to change the password with which Oracle Internet Directory connects to an Oracle database.

### OID Monitor

The Oracle Internet Directory component that initiates, monitors, and terminates the Oracle directory server processes. It also controls the replication server if one is installed.

### one-way function

A function that is easy to compute in one direction but quite difficult to reverse compute, that is, to compute in the opposite direction.

### one-way hash function

A **one-way function** that takes a variable sized input and creates a fixed size output.

### Oracle Call Interface (OCI)

An application programming interface (API) that allows you to create applications that use the native procedures or function calls of a third-generation language to access an Oracle database server and control all phases of SQL statement execution.

### Oracle Directory Manager

A Java-based tool with a graphical user interface for administering Oracle Internet Directory.

### Oracle Internet Directory

A general purpose directory service that enables retrieval of information about dispersed users and network resources. It combines Lightweight Directory Access Protocol (LDAP) Version 3 with the high performance, scalability, robustness, and availability of Oracle8*i*.

### Oracle Wallet Manager

A Java-based application that security administrators use to manage public-key security credentials on clients and servers.

### partition

A unique, non-overlapping directory naming context that is stored on one directory server.

### private key

In public-key cryptography, this key is the secret key. It is primarily used for decryption but is also used for encryption with digital signatures.

### proxy user

A kind of user typically employed in an environment with a middle tier such as a firewall. In such an environment, the end user authenticates to the middle tier. The middle tier then logs into the directory on the end user's behalf, but does so as a proxy user. A proxy user has the privilege to switch identities and, once it has logged into the directory, switches to the end user's identity. It then performs operations on the end user's behalf, using the authorization appropriate to that particular end user.

### public key

In public-key cryptography this key is made public to all, it is primarily used for encryption but can be used for verifying signatures.

### public-key cryptography

Cryptography based on methods involving a public key and a private key.

### public-key encryption

The process in which the sender of a message encrypts the message with the public key of the recipient. Upon delivery, the message is decrypted by the recipient using the recipient's private key.

### public/private key pair

A mathematically related set of two numbers where one is called the private key and the other is called the public key. Public keys are typically made widely available, while private keys are available only to their owners. Data encrypted with a public key can only be decrypted with its associated private key and vice versa. Data encrypted with a public key cannot be decrypted with the same public key.

### referral

See **knowledge reference**.

### relational database

A database is a structured collection of data. In a relational system, data is stored in tables consisting of one or more rows, each containing the same set of columns. Oracle makes it very easy to link the data in multiple tables. This is what makes Oracle a relational database management system, or RDBMS. It stores data in two or more tables and enables you to define relationships between the tables. The link is based on one or more fields common to both tables.

### replica

Each copy of a naming context that is contained within a single server.

### RDN

See **relative distinguished name (RDN).**

### registry entries

Entries containing run-time information associated with invocations of Oracle Internet Directory servers, called **server instance**s. Registry entries are stored in the directory itself, and remain there until the corresponding directory server instance stops.

### relative distinguished name (RDN)

The local, most granular level entry name. It has no other qualifying entry names that would serve to uniquely address the entry. In the example, `cn=Smith,o=acme,c=US`, the RDN is `cn=Smith`.

### replication agreement

A special directory entry that represents the replication relationship among the directory servers in a **directory replication group (DRG)**.

### Root DSE

See **Root Directory Specific Entry**.

### Root Directory Specific Entry

An entry storing operational information about the directory. The information is stored in a number of attributes.

### schema

The collection of **attribute**s, **object class**es, and their corresponding matching rules.

### Secure Hash Algorithm (SHA)

An algorithm that takes a message of less than 264 bits in length and produces a 160-bit message digest. The algorithm is slightly slower than MD5, but the larger message digest makes it more secure against brute-force collision and inversion attacks.

### Secure Socket Layer (SSL)

An industry standard protocol designed by Netscape Communications Corporation for securing network connections. SSL provides authentication, encryption, and data integrity using public key infrastructure (PKI).

### server instance

A discrete invocation of a directory server. Different invocations of a directory server, each started with the same or different configuration set entries and startup flags, are said to be different server instances.

### session key

A key for symmetric-key cryptosystems that is used for the duration of one message or communication session

### SGA

See **System Global Area (SGA)**.

### SHA

See **Secure Hash Algorithm (SHA)**.

### sibling

An entry that has the same parent as one or more other entries.

### simple authentication

The process by which the client identifies itself to the server by means of a DN and a password which are not encrypted when sent over the network. In the simple authentication option, the server verifies that the DN and password sent by the client match the DN and password stored in the directory.

### slave

See **consumer**.

### SLAPD

Standalone LDAP daemon.

### smart knowledge reference

A **knowledge reference** that is returned when the knowledge reference entry is in the scope of the search. It points the user to the server that stores the requested information.

### specific administrative area

Administrative areas control:

- Subschema administration
- Access control administration
- Collective attribute administration

A *specific* administrative area controls one of the above aspects of administration. A specific administrative area is part of an autonomous administrative area.

### sponsor node

In replication, the node that is used to provide initial data to a new node.

### SSL

See **Secure Socket Layer (SSL)**.

### subclass

An object class derived from another object class. The object class from which it is derived is called its **superclass**.

### subschema DN

The list of DIT areas having independent schema definitions.

### subentry

A type of entry containing information applicable to a group of entries in a subtree. The information can be of these types:

- **Access Control Policy Point**s
- Schema rules
- Collective attributes

Subentries are located immediately below the root of an administrative area.

### subordinate reference

A knowledge reference pointing downward in the DIT to a naming context that starts immediately below an entry.

### subtype

An attribute with one or more options, in contrast to that same attribute without the options. For example, a commonName (cn) attribute with American English as an option is a subtype of the commonName (cn) attribute without that option. Conversely, the commonName (cn) attribute without an option is the **supertype** of the same attribute with an option.

### subACLSubentry

A specific type of subentry that contains ACL information.

### subSchemaSubentry

A specific type of **subentry** containing schema information.

### superuser

A special directory administrator who typically has full access to directory information.

### superclass

The object class from which another object class is derived. For example, the object class person is the superclass of the object class organizationalPerson. The latter, namely, organizationalPerson, is a **subclass** of person and **inherits** the attributes contained in person.

### superior reference

A knowledge reference pointing upward to a DSA that holds a naming context higher in the DIT than all the naming contexts held by the referencing DSA.

### supertype

An attribute without options, in contrast to the same attribute with one or more options. For example, the commonName (cn) attribute without an option is the supertype of the same attribute with an option. Conversely, a commonName (cn) attribute with American English as an option is a **subtype** of the commonName (cn) attribute without that option.

### supplier

In replication, the server that holds the master copy of the naming context. It supplies updates from the master copy to the **consumer** server.

### System Global Area (SGA)

A group of shared memory structures that contain data and control information for one Oracle database instance. If multiple users are concurrently connected to the same instance, the data in the instance SGA is shared among the users. Consequently, the SGA is sometimes referred to as the "shared global area."

### system operational attribute

An attribute holding information that pertains to the operation of the directory itself. Some operational information is specified by the directory to control the server, for example, the time stamp for an entry. Other operational information, such as access information, is defined by administrators and is used by the directory program in its processing.

### throughput

The overall rate at which directory operations are being completed by Oracle Internet Directory. This is typically represented as "operations per second".

### trusted certificate

A third party identity that is qualified with a level of trust. The trust is used when an identity is being validated as the entity it claims to be. Typically, the certificate authorities you trust issue user certificates.

### trustpoint

See **trusted certificate**.

### UCS2

Fixed-width 16-bit **Unicode**. Each character occupies 16 bits of storage. The Latin-1 characters are the first 256 code points in this standard, so it can be viewed as a 16-bit extension of Latin-1.

### Unicode

A type of universal character set, a collection of 64K characters encoded in a 16-bit space. It encodes nearly every character in just about every existing character set standard, covering most written scripts used in the world. It is owned and defined by Unicode Inc. Unicode is canonical encoding which means its value can be passed

around in different locales. But it does not guarantee a round-trip conversion between it and every Oracle character set without information loss.

### UNIX Crypt

The UNIX encryption algorithm.

### UTC (Coordinated Universal Time)

The standard time common to every place in the world. Formerly and still widely called Greenwich Mean Time (GMT) and also World Time, UTC nominally reflects the mean solar time along the Earth's prime meridian. UTC is indicated by a z at the end of the value, for example, 200011281010z.

### UTF-8

A variable-width encoding of **UCS2** which uses sequences of 1, 2, or 3 bytes per character. Characters from 0-127 (the 7-bit ASCII characters) are encoded with one byte, characters from 128-2047 require two bytes, and characters from 2048-65535 require three bytes. The Oracle character set name for this is UTF-8 (for the Unicode 2.1 standard). The standard has left room for expansion to support the UCS4 characters with sequences of 4, 5, and 6 bytes per character.

### wallet

An abstraction used to store and manage security credentials for an individual entity. It implements the storage and retrieval of credentials for use with various cryptographic services. A wallet resource locator (WRL) provides all the necessary information to locate the wallet.

### X.509

A popular format from ISO used to sign public keys.

# Index

## Numerics



## A

## B

## C

## D

## E

## F



## G

## H



## I



## J

## K



## L

## M

## N

## O

## P

## Q



## R

## S

## T

## U



## V



## W

## X