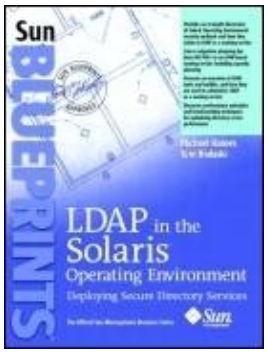

&"87%"
class="v1"
height="17">Table
of Contents
LDAP in the Solaris&"3" class="v2" height="18">By
Michael Haines, Tom Bialaski

Publisher: Prentice Hall PTR
Pub Date: September 17, 2003
ISBN: 0-13-145693-8
Pages: 704

LDAP in the Solaris Operating Environment is a follow-on to the Sun BluePrints book Solaris and LDAP Naming Services, and describes the significant improvements to the Solaris LDAP client and directory server. Deploying the Solaris Secured LDAP Client is covered in detail. This Sun BluePrints book introduces NIS/NIS+ migration tools and techniques to aid in the transition to an LDAP-based naming service. Troubleshooting tips, examples of extending Solaris authentication methods, and examples of extending Solaris authentication methods using the Pluggable Authentication Module (PAM) framework are provided.

[Team LiB]

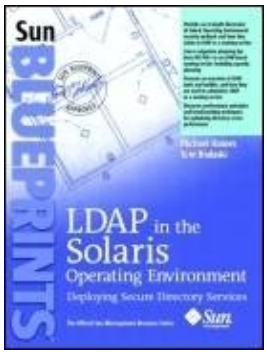

&"87%"
class="v1"
height="17">Table
of Contents
LDAP in the Solaris&"3" class="v2" height="18">By
Michael Haines, Tom Bialaski

Start Reading

Publisher: Prentice Hall PTR
Pub Date: September 17, 2003
ISBN: 0-13-145693-8
Pages: 704

[Team LiB]

Copyright

Prentice Hall PTR offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales. For more information, please contact: U.S. Corporate and Government Sales, 1-800-382-3419, corpsales@pearsontechgroup.com. For sales outside of the U.S., please contact: International Sales, 1-317-581-3793, international@pearsontechgroup.com.

Executive Editor: Gregory G. Doench

Production and Editorial Supervision: Kathleen M. Caren

Cover Design Director: Jerry Votta

Cover Designer: Kavish & Kavish Digital Publishing and Design

Manufacturing Manager: Alexis R. Heydt-Long

Marketing Manager: Debby van Dijk

Sun Microsystems Press:

Publisher: Myrna Rivera

First Printing

Text printed on recycled paper

[Team LiB]

Acknowledgments

This book is made possible through the generous support of many individuals. We sincerely appreciate the time and effort that everyone donated to bring this book to publication.

We would like to give a special thanks to the following people for their invaluable technical expertise and support, which contributed to technical quality and completeness of this book:

Ludovic Poitou Sun&"docText">Neil A. Wilson Sun ONE Product engineer

Jeff T. Moore Solaris™ Naming Services engineer

We would also like to thank the following people for their invaluable technical contribution and support, which also contributed to technical quality and completeness of this book:

Michen Chang Solaris Naming Services engineer

Joep Vesseur Solaris Security Technology engineer

Unai Gaston Caminos Sun ONE Products engineer

Sylvain Duloutre Sun ONE Products engineer

John S. Howard Sun Enterprise™ Engineering engineer

We would like to express our gratitude to those people who contributed to this book through their proficiencies in management and book development:
Many thanks are owed to our families who provided encouragement and understanding when we needed it most. Thank you Jeanne Tringali, Bryony Haines, and Chris Ruble.

Finally, we want to thank the many readers of our first book, Solaris and LDAP Naming Services, Deploying LDAP in the Enterprise. After all, it was those readers who took the time to share their thoughts and experiences with the first book that provided the inspiration and motivation for this book.

[Team LiB]

Preface

LDAP in the Solaris&"docText">This book describes best practices for planning and deploying naming services based on the Lightweight Directory Access Protocol (LDAP). Understanding general LDAP concepts and the specific Solaris implementation is key to successful deployment of resilient enterprise-wide naming services.

This book is a follow-up to the Sun BluePrints book titled Solaris™ and LDAP Naming Service, published in December 2000. The first book introduced LDAP concepts to Solaris system administrators who may not have been familiar with them. It also covered implementation details of the first generation of native LDAP in the Solaris™ Operating Environment (Solaris OE).

Much has changed since the first book was written. The directory server that ships with the Solaris OE has gone through a major revision and several minor ones. The Solaris OE LDAP client software has been significantly enhanced, especially in the area of security. New legacy naming service migration tools have been developed in addition to software that enables co-existence with Microsoft Windows environments.

So much new technology, and so many tools have been developed over the past twoanda-half years, that a simple update to the first book did not make sense. Instead, the content is new. As with the first book, the focus is on how LDAP technology is integrated into the Solaris OE as a naming service, and not a comprehensive book on LDAP concepts and deployments. This book is not meant to replace the Sun product documentation, but rather to complement it by providing expert insight into how the technology works and how best to deploy it. The first book is not a prerequisite for this book.

The Solaris 9 Operating Environment delivers the second phase of Sun's vision for the naming service of the future, and because of the popularity of the Solaris 8 OE, many Solaris 9 OE features have been backported to Solaris 8 OE. New migration tools were included in the first Solaris 9 release and others are included in subsequent updates. The directory server software became integrated in Solaris 9 OE and newer versions are incorporated into Solaris updates.

This book is based primarily on the revisions or software that were available when it was written. Some comparison with older versions is included, so readers who are familiar with those versions can easily understand the differences. This book is based on the following Sun software:

- Solaris 9 4/03 OE
- Solaris 8 OE with Patch 108993-14 (or later version)
- Sun™ ONE Directory Server 5.2 (integrated Solaris OE version)

Many scripts and source code examples are referenced in this book. Rather than including them on a CD-ROM that could quickly become out-of-date, they are posted at http://www.sun.com/solutions/blueprints/tools/index.html. Readers can register, and freely download the examples. See "Obtaining the Downloadable Files for This Book" on page xxvii.

[Team LiB]

Who Should Use This Book

Three types of readers will find the information in this book useful.

- System architects who are responsible for defining enterprise-wide directory and naming service infrastructure.
- System administrators who are tasked with the actual deployment of directory and naming service technology.
- System programmers who must decide on the best way to implement custom features.

[Team LiB]

Before You Read This Book

You should be familiar with the basic administration and maintenance functions of the Solaris OE. You should also have an understanding of standard network protocols and topologies.

Because this book is designed to be useful to people with varying degrees of experience and knowledge about Solaris OE and LDAP technology, your experience and knowledge will determine the path you choose through this book.

[Team LiB]

How This Book Is Organized

This book is organized into the following chapters:

- Chapter 1 "Introducing LDAP in the Solaris Operating Environment" Provides an overview of LDAP-based directory services, the methodologies used to successfully deploy LDAP, and describes terms and concepts commonly used throughout this book.
- Chapter 2 "Assessing Your Needs for Naming Service Transition and Consolidation" Deals with issues of legacy naming services and reasons why you would move to LDAP-based naming services. This chapter presents business reasons for making the transition, and offers tips on migration planning.
- Chapter 3 "Defining Directory Service Security Architecture" Discusses the Solaris OE security model for user authentication and naming service. An example of how to extend the security methods to match your company specific security policies is also provided.
- Chapter 4 "Deploying Solaris OE LDAP Naming Services" Explains methodologies for deploying LDAP as a naming service along with deployment procedures. How to automate the installation and configuration is discussed with step-by-step examples provided.
- Chapter 5 "Migrating Legacy Data to LDAP" Covers migration strategies and the tools that are available for migration. Emphasis is on how to import existing naming service data, and how to configure the directory services to co-exist with legacy naming services.
- Chapter 6 "Management Tools and Toolkits" Provides a survey of tools available from several sources for managing your LDAP naming service data, and provides examples of how to use them effectively. This chapter also describes how to create your own customized tools for managing naming service data.
- Chapter 7 "Performing Administrative Tasks" Presents tricks and tips for administering directory data. The topics covered in this chapter are topics that are not conventionally covered in product documentation.
- Chapter 8 "Selecting Storage for Optimum Directory Server Performance" Describes how to choose the right computer hardware for directory server deployment based on performance characteristics.
- Chapter 9 "Performing Directory Server Benchmarks" Describes the methods and tools used by the Sun Performance Group to characterize the performance of the Sun™ ONE Directory Server software.
- Chapter 10 "Emerging Directory Technologies" Covers important new technologies. These include Directory Service Markup Language (DSML), Sun™ ONE Identity Synchronization for the Windows (ISW) platform and the NIS to LDAP (N2L) transition service.

The following appendices provide supporting material:

- Appendix A, "LDAP Standards Information" Provides references to important documents such as RFCs.
- Appendix B, "LDAP v3 Result Codes" Explains some of the common LDAP error codes that might be returned by your LDAP server.
- Appendix C, "Using snoop with LDAP" Provides information and examples on how to use the `snoop` utility to debug network related LDAP problems.
- Appendix D, "Solaris OE 9 PAM Architecture" Details the `PAM` application programming interface (API) and the PAM service provider interface (SPI). Also included are procedures on how to

effectively write PAM modules when using the Solaris 9 OE.

- The Glossary Provides a list of terms and acronyms used in this book.

[Team LiB]

Obtaining the Downloadable Files for This Book

A variety of files are available for use with this book. The files provide tools and files that are described throughout this book. These downloadable files are available to you free of charge, and by obtaining them, you'll be able to perform many of the best practices described in this book without having to reinvent the code.

We've made every attempt to provide code that is trouble free; however, because every compute environment is different, you should test the code thoroughly before using it in a production environment. As with most freely distributed code, these scripts and tools are provided to you with no support commitment on the part of Sun Microsystems, Inc.

What's Available for Download

TABLE P-1 lists the downloadable files and provides a description of what is included in each file. It is possible that the downloadable files will be updated over time, and those changes might not be listed in this book. In such cases, any changes will be described in a `README` file. Read any `README` files (if present) after unzipping the downloadable file.

Table P-1. Downloadable Files

Downloadable File Name	Description
`ldap-schemas.tar.gz`	Provides two `LDIF` files: • `99nisplusLDAPconfig.ldif` A schema definition for storing NIS+ to LDAP configuration information. See "Configuring the Sun ONE Directory Server Software as a Configuration Server for rpc.nisd" on page 326. • `addSchema.ldif` Schema definition for storing additional NIS+ objects. See "Additional Schema Definitions" on page 329.
`MakeLDIF-1.3.tar.gz`	Provides all the files associated with the `MakeLDIF` utility. `MakeLDIF` is a template-based utility for generating LDIF files. The template file format

	allows for a great deal of flexibility when generating LDIF files. It is very useful for generating data if you want to use SLAMD for benchmarking an LDAP directory server. See "The MakeLDIF Program" on page 480.
`SUNWmakeldif.tar.gz`	Provides the same MakeLDIF files as described above, except these files are delivered in SVR4-package format. To install, use the `pkgadd` command after uncompressing and untarring this downloadable file. See "The MakeLDIF Program" on page 480.
`SUNWmltempl.tar.gz`	Provides template files that give you an automated way of generating your benchmarking data sets using `MakeLDIF`.
`slamd-1.5.1.tar.gz`	Contains the full SLAMD server distribution, including the main SLAMD server, the client application, and the documentation. Refer to the SLAMD Administrator's Guide for documentation on installing the SLAMD server. Also see "Installing the SLAMD Server" on page 510.
`slamd_client-1.5.1.tar.gz`	Contains the SLAMD client application by itself. See "SLAMD Clients" on page 514.
`SOL9-PAM_Modules.class.gz`	Java Installer (WSW) version of the PAM modules: • PAM Compare • PAM Crack • PAM Logon Times Provides the src code, make files (32/64-bit), and man pages. See "PAM Source Code" on page 611.
`SOL9-PAM_Modules.tar.gz`	Provides the same PAM Library modules as described above, except these files are in SVR4-package format. To install, use the `pkgadd` command after uncompressing and untarring the downloadable file. Packages: • `SUNWspamc` PAM Compare Library module • `SUNWspamck` PAM Crack Library module • `SUNWspamlt` PAM Logon Times Library module See "PAM Source Code" on page 611.

`LDAPSubtdel.tar.gz`	The `LDAPsubtdel` program is written in Java and uses classes contained in LDAPJDK 4.1. The tool deletes a specified directory subtree including referrals. See "The LDAPsubtdel Program" on page 394

To Download a File

1. Go to the Sun BluePrints Scripts and Tools download site:

   http://www.sun.com/solutions/blueprints/tools/index.html
2. Select the software that you want to download.

   See TABLE P-1 for downloadable file names.

   Note

   Not all software listed at this site pertains to this book.
3. If you are not already registered at the Sun Download Center, register now.

   You must be registered at the Sun Download Center before you can download the scripts and tools. If this is your first visit, select Register now. You only have to register once, and it's free. Whenever you come back to the Download Center, just enter your user name and password to log in.
4. Log in to the Sun Download Center.
5. Accept the License Agreement.
6. Click on the package listed to download the compressed tar file.

   Perform the download according to the download procedures presented by your browser.
7. Uncompress and untar the file.

```
# gunzip downloadable_file_name
# tar xvf file_name.tar
```


[Team LiB]

Sun BluePrints Program

The mission of the Sun BluePrints Program is to empower Sun's customers with the technical knowledge required to implement reliable, extensible, and secure information systems within the data center using Sun products. This program provides a framework to identify, develop, and distribute best practices information that applies across the Sun product lines. Experts in technical subjects in various areas contribute to the program and focus on the scope and usefulness of the information.

The Sun BluePrints Program includes books, guides, and online articles. Through these vehicles, Sun can provide guidance, installation and implementation experiences, real-life scenarios, and late-breaking technical information.

The monthly electronic magazine, Sun BluePrints OnLine, is located on the Web at http://www.sun.com/blueprints. To be notified about updates to the Sun BluePrints Program, please register on this site.

[Team LiB]

Accessing Sun Documentation Online

The docs.sun.com[sm] web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is http://docs.sun.com/.

[Team LiB]

Typographic Conventions

The following table describes the typographic changes used in this book.

[Team LiB]

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

Table P-2.

Shell	Prompt
C shell prompt	`machine_name%`
C shell superuser prompt	`machine_name"docTableCell" $ valign="top">`
Bourne shell and Korn shell prompt	`#`

[Team LiB]

Chapter 1. Introducing LDAP in the Solaris Operating Environment

Information flow is a critical factor contributing to the success of any large-scale enterprise. This flow of information needs to be rapid, reliable, accessible, and consistent to achieve the maximum benefit. To meet this need, directory services are becoming very popular.

LDAP-based directory services provide a central repository for storing and managing identity profiles, access privileges, and application and network resource information. The information can be used for the authentication and authorization of users, to enable secure access to enterprise and Internet services and applications on a global basis.

Since directory technology is rapidly moving forward with terms and concepts introduced almost monthly, a detailed definition of terms used throughout this book is presented in this chapter with summarized definitions in the glossary for quick reference. Basic LDAP concepts are explained because they are referenced

throughout the book.

This chapter is organized into the following sections:

[Team LiB]

Introduction

A great deal of progress in Solaris OE directory service technology has occurred since Solaris and LDAP Naming Services was published in December of 2000. More enhancements are on the way. This technology area is anything but static, which makes writing a meaningful yet timely book challenging. Because of the number of improvements over the past two years and imminent introduction of even more, we felt this was a good time to publish a follow-up book.

The Solaris 9 Operating Environment delivers the second phase of Sun's vision for the naming service of the future. Included in this release are technologies to securely access an LDAP server from a Solaris OE client and NIS+-to-LDAP migration tools. The LDAP installation and configuration process has also been greatly simplified.

Besides enhancements to Solaris OE naming services, early access to emerging technologies is available. A Directory Service-Markup Language version 2 (DSMLv2) interface is provided with the Sun ONE Directory Server 5.2 software. The Sun&"docText">The Solaris OE authentication framework has been enhanced in Solaris 9 OE to allow incorporation of company-specific security policies. Tools and toolkits based on Java™ technology are available for creating your own customized LDAP management tools.

Drawing on personal experience with early adopters of Solaris OE LDAP naming services, common issues are addressed, as well as the following frequently asked questions:

- What new features does the Solaris 9 and Solaris 8 OE backport client have?
- How do I best deploy these new features?
- What LDAP technologies will be available in the future that might influence my deployment strategy today?
- How do I add functionality to meet my corporate requirements?
- What tools are available for managing LDAP directory data and what are the best practices for their usage?
- How do I integrate Sun's LDAP technology with Active Directory?

These are typical questions asked by IT architects who must define an enterprise-wide LDAP infrastructure, system programmers who need to perform customization, and system administrators who need to develop procedures for deploying and managing LDAP technology.

[Team LiB]

The Big Picture

You are an IT planner for a large multi-national corporation, or perhaps you are a systems manager responsible for a few hundred Solaris OE servers. In any case, you are aware of how critical a naming service

is to the smooth operation of your data center. You may be deploying NIS, NIS+, or your own schema for updating naming service data by propagating text files around the network.

Your current naming service works okay, but you are not totally comfortable with the way it is deployed. Security is a concern and the proliferation of multiple data stores containing user and employee information is creating an administration nightmare. You have also learned that Sun is planning to stop supporting NIS+ in future releases and has similar plans for NIS. LDAP technology is clearly the direction Sun is moving towards as a replacement for NIS and NIS+, but it is not clear what you should be doing about it.

In a nutshell, the purpose of this book is to educate the you about Sun's LDAP implementations so a deployment plan can be established.

In Six Sigma terminology, the steps are:

1. Define
2. Measure
3. Analyze
4. Implement
5. Control

Whether your company is guided by Six Sigma methodology or not, you most likely perform similar activities. This book addresses these five activities in terms of planning and deploying LDAP technology.

Defining the Problem

Quantifying the problem is the first step to developing a solution. In some cases, you might not even be aware that a problem exists. To help in this phase, Chapter 2 "Assessing Your Needs for Naming Service Transition and Consolidation" provides a look at problems typical enterprises face and why they exist.

Measuring the Scope

Recognizing the problem is the first step, and measuring the scope of the problem is the second step. You might be surprised to find out how and where your naming service data is used, and what the authoritative source is. Chapter 2 "Assessing Your Needs for Naming Service Transition and Consolidation" provides a list of common uses and sources of naming service data and can be used as a guide to determining the impact a naming service transition can have.

Analyzing Alternative Solutions

Before you can make an informed deployment decision, you need to know what options are available. There are several chapters that address this activity. Chapter 3 "Defining Directory Service Security Architecture" discusses security options, including how to implement your own security policy. Knowledge about the default behavior of Solaris OE security mechanisms like the Pluggable Authentication Module (PAM) framework is important and included in this book.

Chapter 4 "Deploying Solaris OE LDAP Naming Services" discusses options for deploying a native version of the LDAP naming service client. The term native is used because the client uses LDAP operations to interact with the name service. This is in contrast to the Remote Procedure Call (RPC) interface used by NIS and NIS+ clients. Additional security options covered in Chapter 3 "Defining Directory Service Security Architecture" are specific to the Secured LDAP Client.

Another deployment option is to maintain your current NIS and NIS+ clients, but use an LDAP directory as the back end. The NIS+ migration tool for doing this is covered in Chapter 5 "Migrating Legacy Data to LDAP."

If you have a deployment of Windows 2000 Active Directory servers or Windows NT servers, you may want to visit Chapter 10 "Emerging Directory Technologies" to become familiar with the user account and password synchronization using the Sun ONE Identity Synchronization for the Windows technology.

Implementing

After you have decided on an approach that is right for your enterprise, you need to develop an implementation plan. Chapter 5 "Migrating Legacy Data to LDAP" provides the details on how to implement a native solution. Details on how to implement the NIS+ migration tool are discussed in Chapter 5 "Migrating Legacy Data to LDAP." Another aspect of deployment is converting your current naming service data to LDAP directory data. Techniques for doing this are discussed in Chapter 5 "Migrating Legacy Data to LDAP."

Deciding what hardware is right for your deployment is important. You want to make sure the servers can handle the current load and have headroom for anticipated loads in the future. However, you probably do not have an infinite budget to spend on hardware. Chapter 8 "Selecting Storage for Optimum Directory Server Performance" provides you with guidelines for server sizing and capacity planning.

Controlling

After your initial deployment, the directory data you imported for the LDAP naming service needs to be managed, and client access needs to be controlled. There are many tools and tool kits for doing this. Traditional ones are discussed in Chapter 6 "Management Tools and Toolkits." Emerging technology, namely DSML, for creating tools is covered in Chapter 10 "Emerging Directory Technologies."

[Team LiB]

LDAP Terms and Concepts

Readers who are new to LDAP technology might find some of the terminology confusing. In this section, common terms are defined in the context of LDAP services in the Solaris OE.

Directory Service versus Naming Service

The terms directory service and naming service are often used interchangeably. In a strict sense, a directory service implies that data is stored in a sophisticated structure, while a naming service generally uses a simpler data structure such as a two column table.

LDAP Server versus Directory Server

Often a directory server is mistakenly referred to as an LDAP server because the directory server uses LDAP technology to access information. In this book, the term directory server is used instead of LDAP server, not to be pedantic, but because there is a real difference between the two.

While file servers can be called NFS or SMB servers because of the underlying protocol they use, LDAP is more of an interface that a server uses than a specification of what the server is. A directory server refers to a product or a specific implementation that uses LDAP technology. While there are many directory servers out there, the one most referenced in this book is the Sun ONE Directory Server 5.2 software.

LDAP Models

LDAP is defined by the four models it supports as follows:

- Information Model
- Naming Model
- Functional Model
- Security Model

The following sections describe these models and how they are used in this book.

Information Model

Entries are arranged in a tree-like structure (FIGURE 1-1) called the directory information tree (DIT).

Figure 1-1. Sample Directory Information Tree (DIT)

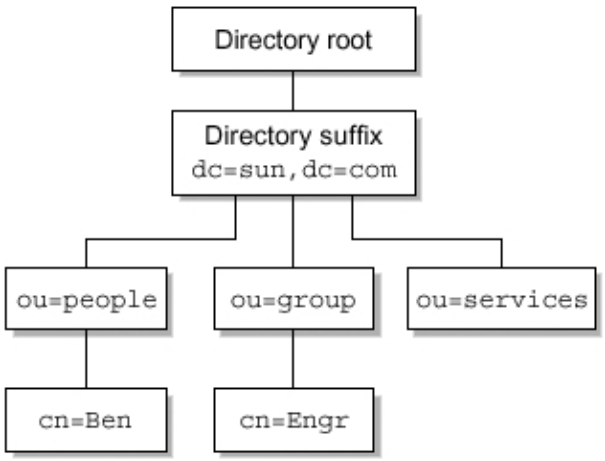

At the top of the DIT is the directory root, which is identified by the server name and port number on which the directory service is running. Multiple instances of the directory service can be running on the same server, with each instance having its own DIT.

Below the directory root is the directory suffix, of which there can be several for each DIT. Suffixes can be expressed as an organization (`o=`) or as an Internet-style domain component (`dc=`). The domain-based format typically mirrors a company's DNS domain address and is expressed as domain component (`dc`) entries.

Located below the suffix are organization unit (`ou`) entries. These entries can be nested, so an ou can contain other organization units. The name chosen for an ou only needs to be unique at the level at which it resides. The same ou can be used in a different portion of the DIT without creating a conflict. An ou entry called `ou=people` is created during the default Directory Server installation. This entry is the default location for storing user account information, but any ou can be used for that purpose.

If there are multiple directory servers in a network, they can be linked by LDAP referrals. A referral is a mechanism that instructs an LDAP client searching the directory to continue the search on another directory server. The referral accomplishes this instruction by passing a Uniform Resource Locator (URL) back to the client. Once the client receives the URL, the client can access the specified directory server.

Naming Model

LDAP is flexible, but at the same time provides enough structure that LDAP clients can access data in any LDAP-compliant directory. For comparison purposes, an LDAP directory is unlike a Solaris OE file system, where a search can always be initiated from the root file system (`/`). Instead, an LDAP directory search begins by specifying one specific entry, such as `dc=blueprints, dc=com`, as a search base. The entry name is specified as a distinguished name (DN), which is a series of relative distinguished names (RDNs).

Each directory server contains a single root directory-specific entry (DSE) which contains basic information about the LDAP server. The DSE contains a list of suffixes supported by the directory server. This entry can be read to determine what suffixes are present, and in turn, those suffixes can be searched.

Functional Model

Clients that need to access data on an LDAP server first perform a bind operation. The act of binding to the directory is analogous to logging in to a system. The bind operation requires, at a minimum, the DN of the user account entry with which the client wishes to bind. If the entry has a password, then, depending on the authentication method, the password is passed along with the DN. Alternatively, the client can perform an anonymous bind, which does not require a particular user name or password.

The bind operation specifies the type of authentication it is attempting to perform. The authentication method the client specifies must be supported by the directory server. The Simple Authentication and Security Level (SASL) interface provides a mechanism for deploying authentication methods that are not directory server specific. The DSE exposes a list of supported SASL mechanisms. The default is Simple authentication, which compares the password sent with the password stored for the specified DN. Because the Simple authentication method sends a clear text password, it is not secure unless combined with Transport Layer Security (TLS), such as SSL. However, other authentication methods (for example, DIGEST-MD5) work securely even without the use of TLS.

If the bind operation is successful, the client is considered authenticated. All subsequent client requests made on the established connection as a result of the bind are performed as the authenticated user. After the LDAP client requests are complete, an unbind operation is performed to release the connection.

Security Model

Access to LDAP entries on the server is protected by the rights established for the authenticated user. The rights can be assigned at the container, object, or attribute level. A portion of the DIT can be assigned stricter (or looser) control than other parts of the DIT. All entries of the same object class type can be assigned the same control. Control can also be established at the attribute level to protect certain information. For example, an employee's password might have restricted access, while other information is available to everyone.

The mechanism used to assign access rights is called the access control instruction (ACI). A single ACI can protect the entire DIT, or several can be used to provide finer-grained protection. When multiple ACIs are created, the ACIs that deny access take precedence. For example,

if access is granted to everyone at the top level of the DIT but denied access to `ou=Contractors`, then the permission set for `ou=Contractors` is enforced.

ACIs are hierarchal. That is, an ACI set at any node of a DIT affects everything below it. Therefore you need to be aware of ACIs that may be set above the node or entry you are protecting.

Directory Objects and Attributes

The structure of a directory entry is defined by the object class to which it belongs. An object class defines a set of attributes that can be stored in a directory entry. LDAP object classes are extensible by creation of a new class that is a child of an existing one. All the attributes defined in the parent class are inherited by the child. The name of an object class must be unique within the directory server and can be registered as a standard LDAP object. These objects are assigned a numeric object identifier (OID) to ensure they will not conflict with another object class.

Attribute names are unique within the directory server and can be contained in more than one object class. The type of data that can be stored in an attribute and the way a search for data is handled is well defined. For example, searching for a string-valued attribute can be case sensitive or insensitive. Attributes can also contain more than one value and can have aliases.

To promote interoperability, a set of standard LDAP object classes and attributes have been defined. Definitions of these ship with most LDAP servers in the form of schema configuration files. If they do not exist on a server, you can add the content of these schema files to the LDAP configuration files.

Directory Schema

The information specified in a directory schema includes the object class name, required and allowed attributes, an OID number, and the allowable syntax. The following table shows the schema definition for the `posixAccount` object class attributes that stores Solaris OE user account information.

Table 1-1. `posixAccount` object Class Schema Definition

Attribute	Description	Syntax
`cn(commonName)`	Common Name of the POSIX account	`DirectoryString` (1-many)
`gidNumber`		`Integer` (single)

	Unique integer identifying group membership	
homePhone	The entry's home phone number	telephoneNumber(single)
uid(userID)	The user's login name	DirectoryString, (single)
uidNumber	An integer uniquely identifying a user	Integer(single)
description	A human-readable description of the object	DirectoryString(single)
gecos	GECOS comment field	DirectoryString(single)
loginShell	Path to the login shell	IA5String (single)
userPassword	Entry's password and encryption method	binary, (single)

In this example, cn is a case-insensitive string that can contain multiple values. The gidNumber and uidNumber are integers, and homePhone is represented by a special data type used for telephone numbers. The loginShell is represented by a case-exact string. Note that the LDAP uid, which is a string, is not the same as the numeric Solaris OE UID, which is represented by the LDAP attribute uidNumber.

Distinguished Names (DN)

A directory entry is identified by its DN, which is similar to a file system path name. Entries are composed of many attributes, some of which are the same as other entries. To distinguish between entries that may have the same values for some attributes, one attribute is usually singled out as being unique. For user account entries defined in the posixAccount object class, the unique attribute is uid.

To prevent duplicate values from being used, the Sun ONE Directory Server software is configured by default to enforce uid attribute uniqueness. Entries that do not have a uid attribute are typically identified by the commonName (cn) attribute, which is available in most object classes, but is not required by all object classes such as organization (o) and organization unit (ou).

The form of a DN is as follows:

attribute=value,container,suffix

There can be multiple containers depending on the DIT topology.

The following is an example of a DN for a user account:

```
uid=cathy,ou=People,dc=example,dc=com
```

The RDN specifies the left-most portion of the DN, which uniquely identifies the entry relative to its parent. For example:

```
uid=cathy
```

In this case, `uid=cathy` has to be unique within the `ou=People` container.

Replication

Replication is the mechanism by which directory data is automatically copied from one directory server to another. Servers that can be updated are called master servers or master replicas and servers that cannot are called read-only replicas. The directory server that provides the data to other servers is called the supplier server. Servers that accept data from suppliers are called consumers. A directory server can be both a consumer and a supplier server. That is, it can receive updates from one server, then pass them on to one or more others.

The Sun ONE 5.2 Directory Server software supports three kinds of replication topologies:

- Single-Master Replication The master copy of directory data is maintained on the supplier server. This server provides updates to consumer servers. All updates are performed on the master or supplier server.
- Multiple-Master Replication Updates can be performed on up to four master servers. The masters keep in sync with each other by propagating updates to each other. In effect, the servers are both suppliers and consumers. A variation of this topology is the floating master. In this configuration, all updates are directed to a single master. If the master server fails, updates are then directed to another one.
- Cascading Replication A server is set up as a hub that supplies updates to several consumers. The hub server receives updates from master servers then passes them on. The purpose of this configuration is to reduce the load on the master server.

Solaris OE LDAP Client

The term LDAP client has two connotations. In the traditional sense, any client that can perform LDAP operations is an LDAP client. This can be an LDAP-enabled browser like Netscape Communicator, or an operating environment that contains commands that are written to the LDAP API. The Solaris OE provides an LDAP API (see `ldap(3LDAP)`), and a number of tools, including `ldapsearch(1)`, `ldapadd(1)`, `ldapdelete(1)`, `ldapmodify(1)`, and `ldapmodrdn(1)` that are built directly on top of the API.

Another use of the term LDAP client is to identify a Solaris OE client that accesses an LDAP directory for its naming service data. The first generation of this client appeared in the Solaris 8 OE. This client is referred to as the Native LDAP Client. The second generation of the client appeared in Solaris 9 OE, and was later backported to Solaris 8 OE. This LDAP client is referred to as the Secured LDAP Client. In general, when the term LDAP client is used in this book, it refers to the Secured LDAP Client implementation.

[Team LiB]

Chapter 2. Assessing Your Needs for Naming Service Transition and Consolidation

This chapter defines the problems businesses face, while making the case for transitioning from a legacy naming service to an LDAP-based directory service. Included are tips on how to identify issues that are commonly encountered in naming service transition and consolidation.