NetWare® 4™ Advanced Administration Student Manual

Course 525

Revision 1.02

Part # 100-001692-003

DPR ..:::: 95

Proprietary Statement

Novell, Inc.
122 E. 1700 S.
Provo, Utah 84606

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of this manual, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes in its content at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any NetWare software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of NetWare software at any time, without obligation to notify any person or entity of such changes.

This Novell Training Manual is published solely to instruct students in the use of Novell networking software. While third-party application software packages are utilized in Novell training courses, this is for demonstration purposes only and shall not constitute an endorsement of any of these software applications. Further, Novell, Inc. does not represent itself as having any particular expertise in these application software packages and any usage by students of the same shall be done at the students' own risk.

Software Piracy

Throughout the world, unauthorized duplication of software is subject to both criminal and civil penalties. For instance, in the U.S., criminal penalties for copyright infringement include fines up to $250,000 and jail terms up to five years, or both. A civil action may result in an injunction, an award of actual damages (including infringer's profits) or statutory damages of up to $100,000 per infringement, and attorneys' fees. If you know of illegal copying of software, call the Novell Anti-Piracy Hotline at 1-800-PIRATES (747-2837). Outside the U.S. call 801-429-7101.

Trademarks

Novell, Inc. has made every effort to supply trademark information about company names, products, and services mentioned in this manual. The following list of trademarks was derived from various sources.

Novell, Inc.

NetWare, the N-Design, and Novell are registered trademarks of Novell, Inc. Btrieve is a registered trademark of Novell, Inc. NetWare Name Service (NNS) is a registered trademark of Novell, Inc. UNIX is a registered trademark of Novell, Inc. WordPerfect is a registered trademark of Novell, Inc.

Certified Novell Instructor (CNI) is a service mark of Novell, Inc. FirstMail is a trademark of Novell, Inc. MHS Services for NetWare 4 is a trademark of Novell, Inc. NetSync is a trademark of Novell, Inc. NetWare 3 is a trademark of Novell, Inc. NetWare 3.12 is a trademark of Novell, Inc. NetWare 4 is a trademark of Novell, Inc. NetWare Client is a trademark of Novell, Inc. NetWare Directory Services (NDS) is a trademark of Novell, Inc. NetWare DOS Requester is a trademark of Novell, Inc. The NetWare Logotype (teeth logo) is a trademark of Novell, Inc. NetWare Message Handling Service (MHS) is a trademark of Novell, Inc. NetWare MultiProtocol Router (MPR) is a trademark of Novell, Inc. NetWire is a registered service mark of Novell, Inc. Packet Burst is a trademark of Novell, Inc. SFT is a trademark of Novell, Inc. Transaction Tracking System and TTS are trademarks of Novell, Inc.

Third Parties

Apple, AppleShare, and AppleTalk are registered trademarks of Apple Computer, Inc. ARCnet is a registered trademark of Datapoint Corporation. DynaText is a registered trademark of Electronic Book Technologies, Inc. IBM is a registered trademark of International Business Machines Corporation. Intel is a registered trademark of Intel Corporation. LocalTalk is a registered trademark of Apple Computer, Inc. Macintosh is a registered trademark of Apple Computer, Inc. Microsoft is a registered trademark of Microsoft Corporation. MS-DOS is a registered trademark of Microsoft Corporation. OS/2 is a registered trademark of International Business Machines Corporation. PostScript is a registered trademark of Adobe Systems Incorporated. Sun Microsystems is a registered trademark of Sun Microsystems, Inc. Xerox is a registered trademark of Xerox Corporation.

Microsoft Windows is a trademark of Microsoft Corporation.

All other products not listed are the property of their respective owners.

Reader Response

Date:	______________________	Number of Pages: ______________
To:	**Course Development Services** **Novell Education**	
Company Name:	**Novell, Inc.**	
FAX Number:	**(801) 429-3900**	

From: ______________________________ Company: ______________________________

Phone Number: ______________________________ FAX Number: ______________________________

We at Novell Education would like to hear your comments or suggestions about this manual. Please fax responses to Course Development Services, (801) 429-3900, or fold and mail this form to:

Course Development Services
Novell Education A-22-1
122 E. 1700 So.
Provo, UT 84606

If you have comments about specific information in this document, please indicate the relevant sections and page numbers.

NetWare® 4™ Advanced Administration — Revision 1.02

Comments ____

Table of Contents

Preface

Course Introduction

Section 1 Server Startup Procedures and Configuration Files

Section 2 Designing and Administering NDS

Section 3 Securing the Directory Tree

Section 4 Partitioning and Replicating NDS and Synchronizing Time

Section 5 Creating a Detailed Design and Troubleshooting NDS

Section 6 Integrating and Managing NetWare 3

Section 7 Configuring NetWare 4 for Diverse Clients

Section 8 Managing Network Services

Section 9 Optimizing the Network and Server

Appendix A

Appendix B

Appendix C

Appendix D

Glossary

Important Certification Information

Novell® guarantees that the certification test based on this course manual will be available for six weeks from the time you acquire the course materials. Thereafter, the test may be replaced at any time by one based on an updated version of the course.

Please arrange to take the test within the six-week period.

The following certification tests are available:

- CNE 050-161
- CNI 050-261

Notes

How to Use This Manual

The *NetWare® 4™ Advanced Administration* manual is designed for use in the *NetWare® 4™ Advanced Administration* course taught by Novell, or Novell Authorized Education Centers. This manual provides training for network administrators responsible for managing a NetWare 4 network operating system.

Keep the following information in mind:

- This course will require three full days. However, the course contains far more material than can be taught in those three days. As a result, your instructor will ask you to read some portions of the course materials outside of class.

You will be tested on all material in the student's manual. You will also be responsible for information in any cross-references in the student manual.

- We strongly recommend that you complete the prerequisite courses in the CNE program (or gain equivalent experience) prior to taking this course. If you need help on the basics in any area, such as printing, discuss this with your instructor as soon as possible.
- Each section contains several objectives. You should be able to complete each objective. The CNE/CNI tests are based on these objectives. Novell tests are moving toward being "performance-based" and "job-task-centered." The test will evaluate what you can do as well as what you can remember.
- Share relevant experiences with the class.
- You should practice the skills you learn in this class. The sooner you practice the skills, the better you will remember what you learned.

Product Documentation

The manual is written in workbook format with brief concept explanations and extensive space for taking notes. It is designed to accompany the NetWare product documentation, which contains the technical detail pertaining to concepts presented in the class. Where appropriate, this manual refers you to specific information in the documentation. Documentation references are at the bottom of the page.

Exercises

The *NetWare® 4™ Advanced Administration* course is designed to provide extensive hands-on training with NetWare 4. Most sections present hands-on exercises for you to perform. Some sections also contain other practice exercises.

Section Objectives

Objectives are listed at the beginning of each section. Use these at the beginning of each section as an overview of the material to be covered. Use them after completing a section to verify that you understood the most important concepts and issues within that section.

Icons

Provides more specific information on concepts or procedures, such as sample cases or command line entries.

Emphasizes critical information, such as warnings or special instructions.

Provides additional information or explanation of possible results (nice, not critical).

Lists product documentation to be used as a resource for a course topic or activity.

Recommends useful ideas or actions to apply for a given topic.

Notes

NetWare® 4™ Advanced Administration Course

Questions and Answers

This form will help answer any questions that you may have concerning the *NetWare® 4™ Advanced Administration* course and NetWare products.

If you have specific questions that are not answered in class, or that occur to you later, please fill out this form and give it to the instructor at any time. The instructor will answer your questions in class if time allows; otherwise, you will receive the answers by mail.

Name: ________________________________

Company: ________________________________

Address: ________________________________

Date: ________________________________

Questions: ________________________________

1.

2.

3.

4.

5.

6.

7.

8.

9.

10.

Novell's Anti-Piracy Program

As a future Certified Novell Engineer, Administrator, or Instructor, you will have the opportunity to work with challenging issues on a daily basis. The proper use and licensing of NetWare can be a difficult job for a networking professional when working with or for others who may not understand the importance of strict license compliance. Piracy of NetWare, along with the software of many other manufacturers, is a multi-billion dollar problem world-wide. The Business Software Alliance (BSA) calculated that in 1993, 35 percent of all software used in the United States was pirated. This translates into a $1.9 billion annual loss for software developers.

When an act of piracy occurs, the following people may be affected:

- The victim, who has paid for original software and has not received original disks, manuals, and registration materials
- The honest reseller, who cannot compete with pirates and loses business
- The author or creator of the software, who loses money and whose reputation suffers from poorly installed, pirated software
- The Novell-educated service provider, who loses clients because users who fear being reported to Novell try to find support from noncertified providers who will not ask any questions

What Is Piracy?

Software piracy occurs when software is installed or distributed in more locations than the license allows. A *victim of piracy* is someone who has paid for software, but has unknowingly received an illegal copy. Novell wants to help victims, not complicate their problems.

What Do I Look For?

The following may indicate software piracy:

- Various users running the same serial number.
- Lack of original documentation. This often indicates piracy, but is *not always* the case.
- Any sign that indicates copied software is being used.

Although most of Novell's efforts focus on reseller user software piracy, we also deal with trademark/trade name abuse, which includes the following:

- An unauthorized reseller using Novell's trademarks or trade names in an advertisement or display
- An unauthorized reseller claiming to be authorized
- An unauthorized claim of "NAEC" or "CNE" certification

How Do I Report Piracy?

When you suspect an illegal installation, report the suspected piracy to Novell by calling the **toll-free** number (1-800-PIRATES). Your call will be kept confidential. **Do not** contact your distributor or the Novell sales office. Call the Piracy hotline to find out where to send the information.

In the unlikely event your employer knowingly runs illegal software or asks you to make an illegal installation, explain the restrictions of the software license and the potential for legal action by Novell. Encourage your employer to purchase legal versions that can be registered and supported. By not reporting instances of pirated software or by participating in the infringement of Novell copyrights, CNEs, CNIs, and CNAs are in danger of losing their certification.

How Do Honest Service Providers Benefit?

By helping these users, Novell and the service provider will gain an appreciative customer who is now eligible for upgrades and support. You, as the service provider, will be in a position to provide those services to the new customer. You will also be reducing the unfair competition in your area.

How Are Cases Resolved?

When piracy of NetWare is established, Novell will do one of the following:

- Send a letter seeking resolution.
- Conduct an audit to determine the scope of the problem.
- Start civil litigation.
- Start criminal litigation.

Whenever possible, Novell will also issue a press release. Novell believes that increased public awareness will decrease piracy and increase legitimate sales and service.

Novell provides a toll-free hotline (1-800-PIRATES) to answer your questions or concerns. All calls will be kept confidential. It is Novell's policy to reveal neither the source of the information nor the identify of the complainant. To contact this hotline from outside of the United States, call 801-429-7101.

Report all instances of known or suspected piracy
1-800-PIRATES (747-2837). 801-429-7101 (outside of United States)

Notes

Course Introduction

Introduction

This section provides an overview to the course, the materials, and other administrative activities. Student reference materials are discussed and NetWare product documentation is also presented.

Objectives

At the end of this section, you will be able to perform the following tasks:

1. Explain the sections in the course and the layout of the materials.
2. Explain the materials used in the course, which include the student manual, documentation, and software.
3. Identify student manual references to NetWare® product documentation.
4. Review tasks for using Novell® online documentation.

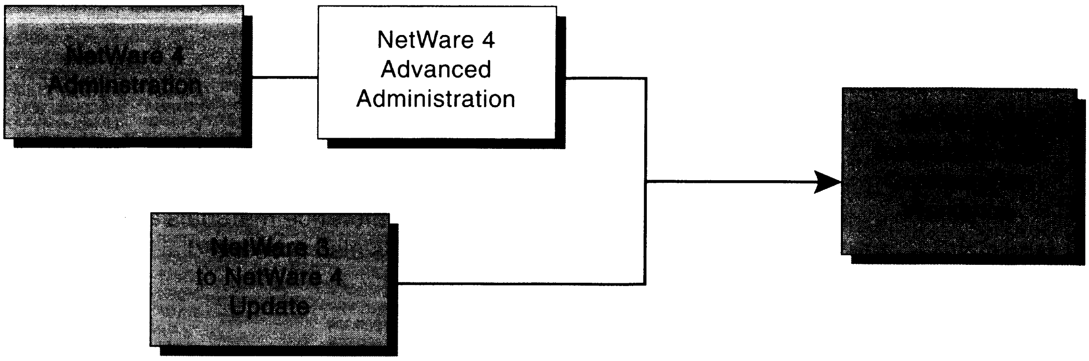
NetWare 4 Administration
NetWare 4 Advanced Administration
NetWare 3 to NetWare 4 Update

Figure Intro-1

NetWare 4 Curriculum

Several courses in the NetWare 4™ curriculum teach network administrator responsibilities. This course is one of four courses available. Make sure you fit the profile for this course. This course has been designed for NetWare 4 administrators who have attended the NetWare 4 Administration course. If you are working with NetWare for the first time, the NetWare 4 Administration course would be more appropriate.

Course	Responsibilities	Description
NetWare 4 Administration	Manage, Protect, Back Up, Organize File System	Set up and manage basic network services: create users, set up file systems, set up security, and manage printing.
NetWare 4 Advanced Administration	Manage, Tune, Set Up, Organize	Perform more complex tasks: plan directory structure, tune performance, and troubleshoot NetWare 4.
NetWare 3 to NetWare 4 Update	Manage, Tune, Set Up, Organize the Directory	Manage an existing network: create users, set up security, and manage printing. Perform more complex tasks: plan directory structure, tune performance, and troubleshoot NetWare 4.
NetWare 4 Installation and Configuration Workshop	Install, Upgrade, and Migrate NetWare 4 Configure the Server and NDS™	Install NetWare 4, determine appropriate settings for installation, and create and manage NetWare Directory Services™ (NDS).

Table Intro-1: NetWare 4 Course Descriptions

Course Map

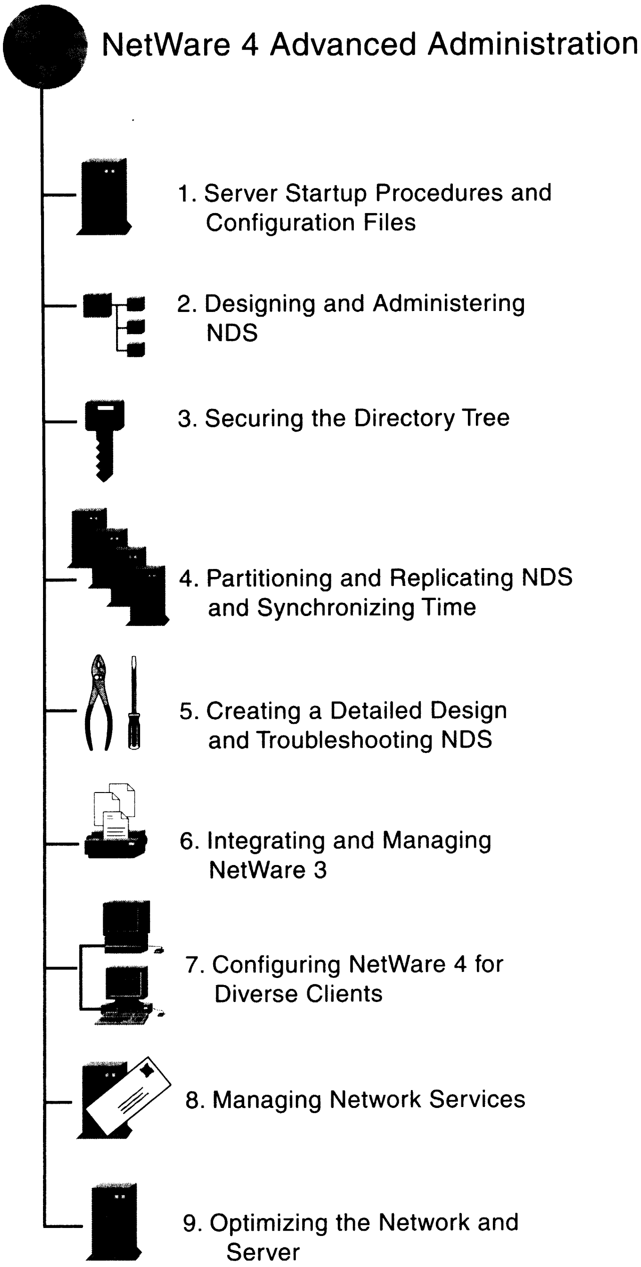

NetWare 4 Advanced Administration

1. Server Startup Procedures and Configuration Files
2. Designing and Administering NDS
3. Securing the Directory Tree
4. Partitioning and Replicating NDS and Synchronizing Time
5. Creating a Detailed Design and Troubleshooting NDS
6. Integrating and Managing NetWare 3
7. Configuring NetWare 4 for Diverse Clients
8. Managing Network Services
9. Optimizing the Network and Server

Figure Intro-2: NetWare 4 Advanced Administration Course Map

Course Goal

This course is designed to provide students with the knowledge and skills to design, configure, and administer a complex network. Students who complete this course will be able to accomplish advanced network tasks on a NetWare 4 network. The course is designed to provide students who have completed the NetWare 4 Administration class with an advanced skill set and ability to handle more challenging network situations than were presented in the basic administration course.

Course Objectives

At the end of this course, you will be able to perform the following tasks:

1. Review tasks for using Novell online documentation. (Introduction)
2. Perform server startup procedures and maintain the server's configuration files. (Section 1)
3. Identify the process for creating a design for the Directory tree; this process includes determining the Directory structure, naming conventions, and implementation method. You will also learn how to create several users at one time. (Section 2)
4. Secure the Directory tree by applying access control with trustee assignments to containers, implementing Inherited Rights Filters (IRFs), and creating NDS security strategies based on guidelines. (Section 3)
5. Determine strategies and guidelines for partitioning, replicating, and distributing the Directory database. (Section 4)
6. Configure servers to synchronize time on a NetWare 4 network. (Section 4)
7. Given scenarios, create a detailed design and troubleshoot NDS replica loss, from down time and inconsistencies through synchronization. (Section 5)
8. Integrate and manage NetWare 3 resources with NetWare 4 utilities. (Section 6)
9. Install and configure the NetWare DOS Requester™ and the server for diverse clients that include OS/2*, Macintosh*, and UNIX®. (Section 7)

10. Configure multiple protocols and messaging with the NetWare 4 server. (Section 8)

11. Identify options for configuring the NetWare 4 server with languages, complex file systems, and routing. (Section 8)

12. Given scenarios, optimize the NetWare 4 server through monitoring server statistics and taking appropriate actions with SET commands and other server utilities. (Section 9)

Reference Materials

This course uses the following materials to provide you with training information:

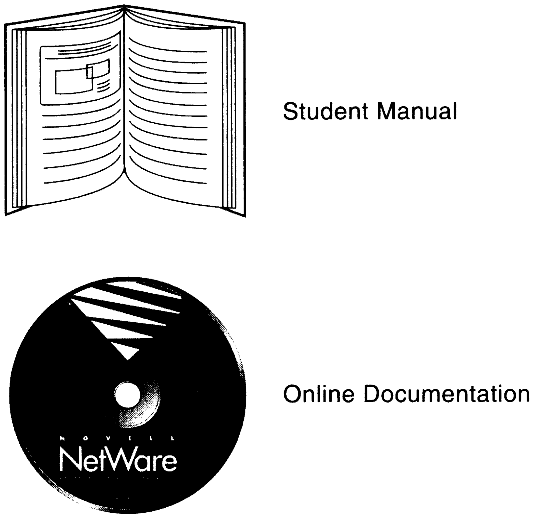

Figure Intro-3: Student Reference Materials

Student Manual

The student manual provides concept organization and presentation, class exercises, and space for notes. It does not contain extensive product information. See "How to Use This Manual" in the Preface for information on the following subjects:

- Product documentation
- Exercises
- Section objectives
- User Comments

Online Information

NetWare 4 provides access to several forms of online information. It comes in the following forms:

- Command line help
- Menu utility help
- F1 Help for the graphical utilities
- Novell online documentation

Discussions of and exercises using these features are introduced later in the course as needed.

Novell Online Documentation

NetWare 4 provides electronic versions of all its product documentation. The DynaText utility allows you to choose a manual, browse its contents, search for topics, and print the content.

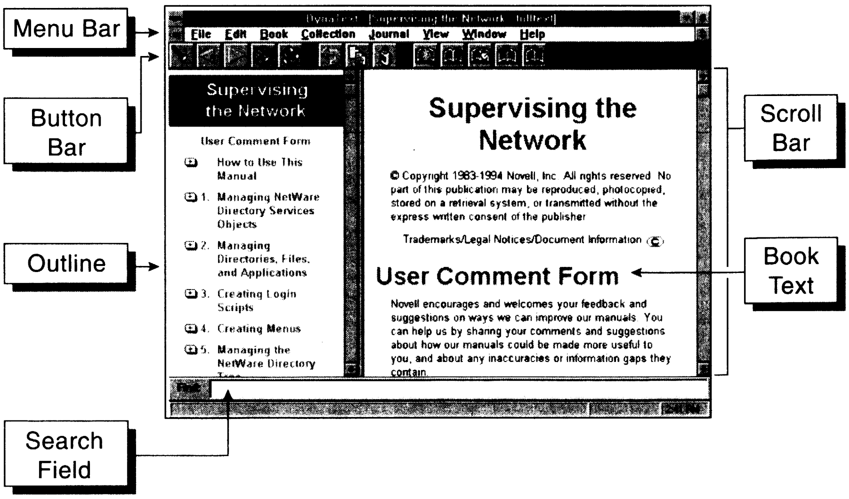
Menu Bar
Button Bar
Outline
Search Field
Scroll Bar
Book Text
DynaText - [Supervising the Network - fulltext]
File Edit Book Collection Journal View Window Help
Supervising the Network
User Comment Form
How to Use This Manual
1. Managing NetWare Directory Services Objects
2. Managing Directories, Files, and Applications
3. Creating Login Scripts
4. Creating Menus
5. Managing the NetWare Directory Tree
Find

Supervising the Network

© Copyright 1983-1994 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Trademarks/Legal Notices/Document Information

User Comment Form

Novell encourages and welcomes your feedback and suggestions on ways we can improve our manuals. You can help us by sharing your comments and suggestions about how our manuals could be made more useful to you, and about any inaccuracies or information gaps they contain.

Figure Intro-4: Novell Online Documentation

The main areas of the interface include the following:

- Outline
- Book Text
- Search Field
- Button Bar
- Menu Bar
- Scroll Bar

The NetWare 4 product documentation provides more detailed reference information than the student manual. It comes in two forms: printed and online. References to the product documentation can be found throughout this student manual at the bottom of the page. The references are identified by the following icon:

📖

📖 For more information on using DynaText and the Novell online documentation, see *Installing and Using Novell Online Documentation.*

Summary

This course is the NetWare 4 Advanced Administration course. Tasks include administering NetWare 4 and designing NDS. Many sources of reference material are provided with this course.

SECTION 1 Server Startup Procedures and Configuration Files

Introduction

In this section, you will gain an understanding of the NetWare 4 server startup process, including some of the options for modifying that process. You will also create server batch files to automate server tasks.

Objectives

At the end of this section, you will be able to do the following:

1. Identify and describe the server components.
2. Perform a server startup procedure.
3. Identify and describe server configuration files.
4. Identify commands and options used to customize the appropriate server configuration file.
5. Select the appropriate utility and edit the server configuration files.
6. Create server batch files that perform specific tasks, such as remotely rebooting the server.

Server Components Overview

The following subsections summarize the hardware and software components and startup procedures for the NetWare 4 server.

Hardware Components

NetWare 4 requires the following hardware components:

- Processor - 386 or higher

  NetWare 4 requires at least an Intel® 386 processor. The more services provided, the faster the processor that is required.

- RAM

  NetWare 4 requires at least 8 MB of RAM. For details on calculating the appropriate RAM for your environment, refer to Appendix A of the *Supervising the Network* manual or to Course 804, *NetWare 4 Installation and Configuration Workshop.*

- Disk storage

  NetWare 3 and NetWare 4 servers use the same disk system environment: a DOS partition for SERVER.EXE and a SYS: volume on the NetWare partition (see Figure 1-1). You can create additional volumes using the same hard drive, a CD-ROM, or other additional storage devices.

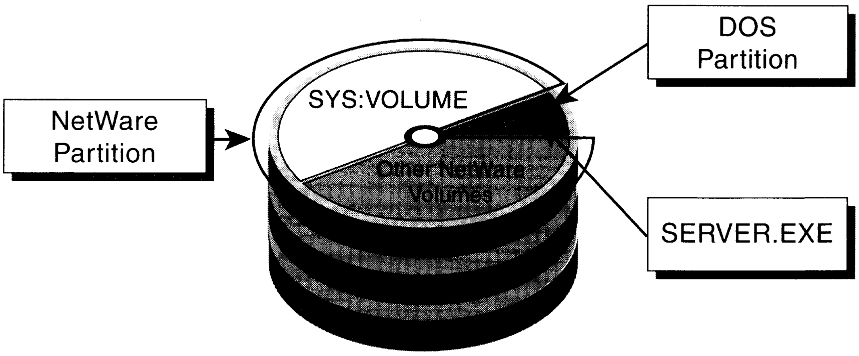

Figure 1-1: NetWare Disk Storage

- Network board and cabling

  The server requires a network board and cabling (or an equivalent wireless connection) to communicate with clients and other devices. The server supports multiple network boards simultaneously and allows communications to occur between the networks. In this capacity, the server provides a routing function.

Software Components

The core software components of a NetWare server include the following: SERVER.EXE, a disk driver, and a LAN driver.

- SERVER.EXE is the primary NetWare operating system file. Once in RAM, SERVER.EXE takes over the operation of the computer from DOS.
- A disk driver allows the server to communicate with the hard disk.
- A LAN driver allows the server to communicate with the network.

Server Identification

NetWare networks use numbering to uniquely identify the server on the network. The numbering identifies the server, network board, and cable segment, as shown in Figure 1-2.

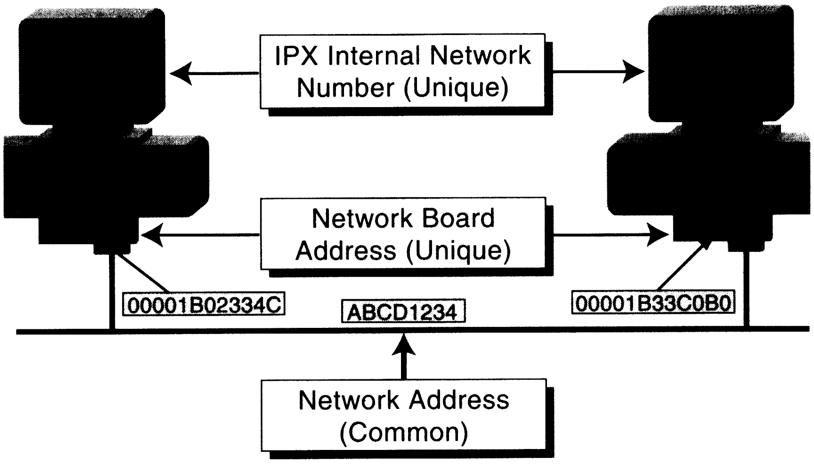
IPX Internal Network Number (Unique)
Network Board Address (Unique)
00001B02334C
ABCD1234
00001B33C0B0
Network Address (Common)

Figure 1-2: Server Network Numbering

Startup Procedures

The NetWare 4 server startup procedure includes the following steps, which may be automated in batch files.

1. Execute SERVER.EXE.
2. Load the disk driver.
3. Provide the server name and the internal IPX network number.
4. Load the LAN driver.
5. Bind the LAN driver to the network cable segment.

Exercise 1-1: Executing the Server Startup Process

Activity: Review the NetWare 4 server startup.

Procedure: Identify the files on the DOS partition, boot the server, and view server information with the server utilities.

1. View the server files on the DOS partition.

   a. At the server DOS prompt, change to the **NWSERVER** directory.

   b. Type

      DIR <Enter>

      Notice all the .NLM and .DSK files.

   c. Type

      DIR *.EXE <Enter>

      How many .EXE files do you see?

> Before you perform Step 2, make sure no memory managers have been loaded by the CONFIG.SYS file.

2. View the server startup process by typing

   SERVER <Enter>

   View the messages as the server starts up.

3. Check configuration and network information at the console prompt.

   a. To view configuration information at the console prompt, type

      CONFIG <Enter>

      If you are using Ethernet, notice that the default frame type for your LAN driver is Ethernet 802.2.

   b. To see other servers in the network, type

      DISPLAY SERVERS <Enter>

      Look for your server name in the list. The DISPLAY SERVERS command shows Directory trees and servers. Each server may appear several times, once for each service that it advertises. Common services are NetWare Directory Services (NDS) and file services.

4. View server statistics with MONITOR.NLM.

   a. Type

      LOAD MONITOR <Enter>

   b. Notice what happens when you press **<Tab>**.

   c. Select other MONITOR options of interest to you.

   d. Exit MONITOR.

5. View the configuration files and server setup with INSTALL.NLM.

   a. Type

      LOAD INSTALL <Enter>

   b. Select **Volume options**.

   c. Press <Enter> to view the information on the SYS: volume.

      NetWare 4 compression works on a file-by-file basis. Files are compressed if they are not used for a specified amount of time.

      Block suballocation divides any partially used disk block into 512-byte suballocation blocks. These suballocation blocks are used to share the remainder of the block with leftover fragments of other files.

      These two features of NetWare 4 provide efficient use of your file system.

   d. Escape out of this view and select **NCF file options**.

   e. View the existing **AUTOEXEC.NCF** file.

      Notice the commands listed after the SET commands. These commands are the same in NetWare 3 and NetWare 4.

   f. Exit the INSTALL utility.

You have now completed the exercise.

Server Configuration Files

Configuration files are text files created to specify actions that should occur when a device is booting. A secondary benefit of configuration files is that these files automate the startup procedures.

DOS devices use both CONFIG.SYS and AUTOEXEC.BAT to create the DOS environment and automate the startup procedures of the device.

Server configuration files are similar to DOS configuration files; however, they include commands that are specific to the NetWare server.

A server may also use AUTOEXEC.BAT to automatically execute SERVER.EXE. Once loaded, the NetWare operating system searches for its configuration files. The server's configuration files are listed below:

- STARTUP.NCF
- AUTOEXEC.NCF

The order in which these files are used is shown in Figure 1-3.

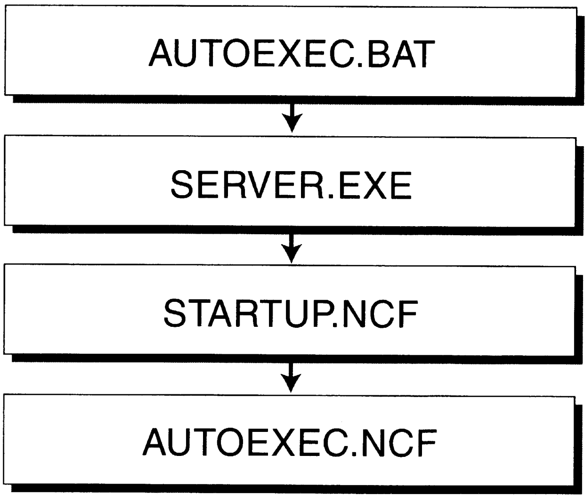

Figure 1-3: Execution Order of Server

For more information about AUTOEXEC.NCF and STARTUP.NCF, see Chapter 3, "Custom Installation" in *NetWare 4 Installation.*

STARTUP.NCF

STARTUP.NCF is a server boot file that contains the commands to load disk drivers. It contains the configuration parameters to load the disk driver and may contain other parameters to configure the server. These parameters are part of the SET command and cannot be used at the server console prompt.

Other commands may be added to accommodate filenames other than DOS in the server's file system. These commands are called name space commands. A sample NetWare server STARTUP.NCF file might contain the following commands:

```
LOAD ISADISK PORT=1F0 INT=E
SET MAXIMUM SUBDIRECTORY TREE DEPTH=10
```

STARTUP.NCF can be generated automatically by the INSTALL utility during installation and stored in the server's boot directory. It can be edited through INSTALL.

To boot with an alternate STARTUP file (for maintenance purposes), use the following syntax:

SERVER -S \path\filename

To boot without a STARTUP.NCF file (for maintenance or troubleshooting), use the following syntax:

SERVER -NS

AUTOEXEC.NCF

The AUTOEXEC.NCF file, located in SYS:SYSTEM, guides the file server during the boot process. It contains the following information:

- Time zone information
- Bindery context information
- File server name
- Internal network number
- Calls to other .NCF files
- Other LOAD commands for NLMs (except those NLMs which belong in STARTUP.NCF)
- Other console commands (including SET commands)
- Comment lines (marked with a semicolon in the example below)

A sample AUTOEXEC.NCF file might contain the following commands:

```
;ECORP'S AUTOEXEC.NCF FILE, LAST MODIFIED 1-11-95 BY BOB CUTLER
SET TIME ZONE = MST7MDT
SET DAYLIGHT SAVINGS TIME OFFSET = 1:00:00
SET START OF DAYLIGHT SAVINGS TIME = (APRIL SUNDAY FIRST 2:00:00 AM)
SET END OF DAYLIGHT SAVINGS TIME = (OCTOBER SUNDAY LAST 2:00:00 AM)
SET DEFAULT TIME SERVER TYPE = SINGLE
SET BINDERY CONTEXT = OU=CORP.O=EMA
FILE SERVER NAME ECORP
IPX INTERNAL NET 1011532
LOAD NE2000 INT=3 PORT=300 FRAME=ETHERNET_802.2 NAME=NE2000_1_E82
BIND IPX NE2000_1_E82 NET=010107D0
LOAD REMOTE CONSOLE
LOAD RSPX
LOAD MONITOR
SET SOUND BELL FOR ALERTS = OFF
```

This file is created by INSTALL during installation, but you can also create or edit it with SYSCON.

To boot without an AUTOEXEC.NCF file (for maintenance or troubleshooting), use the following syntax:

SERVER -NA

Customizing the Server Configuration Files

The configuration files can be customized to meet your needs for management and server performance. NLMs, console commands, and SET parameters can be added to provide the services you need.

The following utilities can be used to edit STARTUP.NCF and AUTOEXEC.NCF:

- INSTALL.NLM
- EDIT.NLM

For more information, see each command or utility in *Utilities Reference.*

NLMs and Console Commands

The following console commands and NLMs may be useful in customizing your configuration files.

Command	Syntax	Description
REMOTE RSPX	LOAD REMOTE *password* LOAD RSPX	Load remote console support
VOLUMES	VOLUMES	Display mounted volumes
CONFIG	CONFIG	Display server configuration
MODULES	MODULES	Display NLMs that are loaded
SEARCH	SEARCH ADD *path*	Add location to search for NLMs
REMOVE DOS	REMOVE DOS	Remove COMMAND.COM from server RAM; reboot the server when EXIT is executed
SECURE CONSOLE	SECURE CONSOLE	Remove DOS from RAM; require cold boot when EXIT is executed
MONITOR	LOAD MONITOR	Load monitoring utility

Table 1-1: Common NLMs and Console Commands

Any of these commands and NLMs can be loaded from the console prompt. If you always want them available when the server boots, place the commands in the AUTOEXEC.NCF file.

SET Parameters

The default server configuration parameters in NetWare 4 have been set to maximize network performance in most situations. You can view or change these parameters using the SET command. The following parameter categories appear on the SET console screen:

- Communications
- Memory
- File caching
- Directory caching
- File system
- Locks
- Transaction tracking
- Disk
- Time
- NCP
- Miscellaneous
- Error handling
- Directory services

To view the SET console screen, enter SET at the server console.

To change a parameter's setting, type the following command:

SET *parameter* = *value* <Enter>

This command format can be used to customize the server configuration files. Certain parameters must appear in the STARTUP.NCF file; others are added to the AUTOEXEC.NCF file.

Another way to change SET parameters is with the SERVMAN console utility. This utility provides a menu interface to change parameters and save those changes automatically to the configuration files.

For more information on SET commands, see "SET" in *Utilities Reference*.

Some useful parameters are outlined in Table 1-2 below. The default values are displayed with the parameter.

SET Parameter	Values	Notes
Maximum subdirectory tree depth=25	10 to 100	Restricts the depth of directories allowed in the directory structure of a NetWare volume
Reply to get nearest server=ON	ON, OFF	Allows a server to respond to a "Get Nearest Server" request from a client or prevents a server from responding to the request
Display lost interrupt alerts=OFF	ON, OFF	
Allow unencrypted passwords=OFF	ON, OFF	

Table 1-2: Sample SET Parameters

Most SET parameters can be executed at the server console prompt, but the setting will last only until it is reset or until you bring down the server. To make the settings permanent, add them to the appropriate configuration files.

Exercise 1-2: Customizing the Server AUTOEXEC.NCF File

Activity: Customize the server AUTOEXEC.NCF file by adding remote management NLMs and changing SET parameters.

Procedure: Complete the following tasks.

1. At the server console screen, type

   LOAD INSTALL <Enter>

2. Select **System Options** from the Installation Options menu.
3. Select **Edit AUTOEXEC.NCF File** from the Available System Options menu.
4. Add the commands to load the remote console software to AUTOEXEC.NCF. Save the changes. The remote console commands are

   LOAD REMOTE *password*
   LOAD RSPX

5. Exit INSTALL.
6. At the server console, bring down the server and exit to DOS.

   a. Type

      DOWN <Enter>

   b. Type

      EXIT <Enter>

7. Restart your server, watch the console messages, and respond to the prompts. Type

   SERVER <Enter>

8. To make sure the REMOTE and RSPX modules loaded, type

   MODULES <Enter>

Creating Server Batch Files

mnt sk flr

You can create batch files to automate procedures that you execute at the server console. These files can be executed from the server console prompt or called from other batch files, such as AUTOEXEC.NCF.

Batch files executed from the server console require the .NCF extension.

You can create the batch files with a text editor and place them in the SYS:SYSTEM directory, or you can create and edit them at the server console using EDIT.NLM.

Use the following syntax to load EDIT:

> **LOAD EDIT** *path\filename*

When you use EDIT, you must specify the location of the file and the filename. You must enter a file extension if you want one; EDIT will not create an extension for you.

EDIT.NLM can also be used to edit text files (such as AUTOEXEC.BAT) on the DOS partition of the server, provided you have not removed DOS from server memory.

Exercise 1-3: Creating Server Batch Files

Activity: Create a server AUTOEXEC.BAT file and a batch file that together can remotely reboot the server.

Procedure: Using RCONSOLE.EXE, complete the steps below to create the following files.

Start a remote console session with your server.

1. Reboot your workstation.
2. Log in to your server. Type

   LOGIN E*nnn*/admin.EMA*nnn* <Enter>

3. Establish a remote console session with your server.
   a. Type

      RCONSOLE <Enter>

   b. Select **SPX**.
   c. Select your server name.
   d. Type in the password you used with the LOAD REMOTE command in the server's AUTOEXEC.NCF file.

Create the server's AUTOEXEC.BAT file to load SERVER.EXE.

4. Type

   LOAD EDIT C:\AUTOEXEC.BAT <Enter>

5. Enter the following commands in the file:

```
CD \NWSERVER
SERVER
```

6. Exit and save the file.

Create a batch file (named REBOOT.NCF) to reboot the server.

7. Type

   LOAD EDIT SYS:SYSTEM\REBOOT.NCF <Enter>

8. Enter the following commands in the file:

   ```
   REMOVE DOS
   DOWN
   EXIT
   ```

9. Exit and save the file.

Reboot the server.

10. Type

    REBOOT<Enter>

    You will see a message that the server has open files.

11. Respond to the "Down server?" prompt with **Y**.

12. Observe the server rebooting and autoloading SERVER.EXE.

Summary

This section lays a foundation for understanding a NetWare server. The information presented in this section will be used throughout the rest of this course. This section discussed the basic hardware and software components needed to create a NetWare 4 server. It discussed the server startup procedure, which initializes SERVER.EXE, disk drivers, and LAN drivers. Finally, it discussed the creation of batch files, such as STARTUP.NCF and AUTOEXEC.NCF, which are used to automate the server startup procedure.

Notes

SECTION 2 Designing and Administering NDS

Introduction

In this section, you will learn basic Directory tree design concepts. By designing your Directory tree before creating it, you can plan how users will access the Directory. Although you can use NDS™ utilities to change your Directory tree, you will save time and effort by following the guidelines presented in this section.

You will also learn how to implement tree designs for various approaches. For example, you will learn how to set up a Directory tree for a small department and then merge the tree into a larger organization. You will also learn how to create a tree for an organization and add upgraded servers to the tree.

Administering a Directory tree requires some navigation skills. This section will teach you how to move around in the tree and specify objects in the way NDS requires. You will also learn how to set up users so that they can access network resources more efficiently and benefit from the advantages of NDS.

Objectives

After you complete this section, you will be able to do the following:

1. Identify the process for designing the Directory tree, which includes determining the Directory structure, naming conventions, and implementation method.

2. Create a structural design for an organizational, divisional, or departmental Directory tree implementation.

3. Merge small Directory trees in a departmental-to-organizational implementation.

4. Navigate the contexts of the newly created Directory tree using CX and NetWare Administrator.

5. Demonstrate correct use of object names in utilities.

6. Create users with UIMPORT.

Exercise 2-1: NDS Review

This will be a group exercise. Your instructor will provide directions.

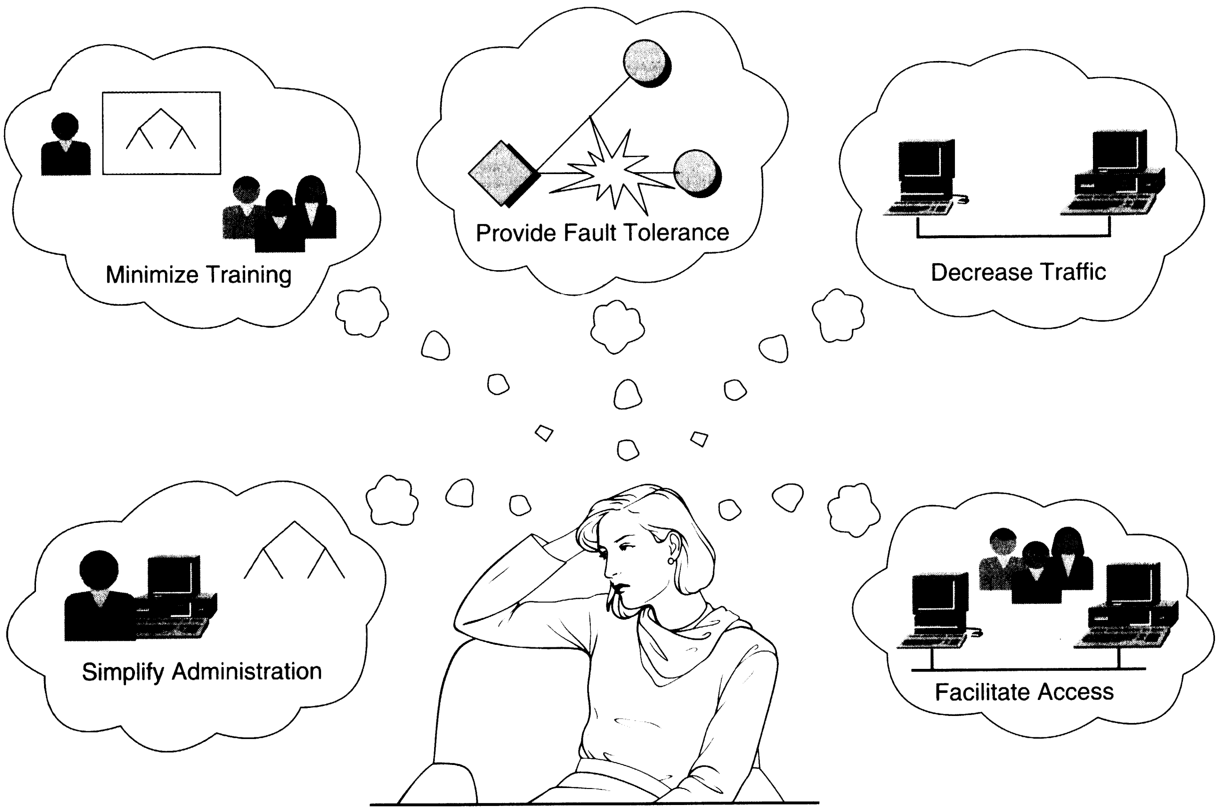

Figure 2-1

Benefits of an Effective Directory Tree Design

Planning your Directory tree effectively allows you to do the following:

- Simplify network administration and maintenance.
- Minimize the impact on users and reduce the need for training.
- Enable users to access network resources easily.
- Provide NDS fault tolerance for the network.
- Decrease unnecessary network traffic.

There is no one “right” way to design an NDS tree. You need to design a tree that meets your organization’s needs. You can modify the Directory tree later if your needs change.

Identifying Components of Directory Design

You can take several approaches in determining how you will design and implement NetWare 4. The approach of this course is divided into two phases:

- Structural design
- Detailed design

Structural Design

The structural design focuses on the structure of the Directory tree and the process for implementing your structure. Structural design is covered in this section. The steps in the structural design are shown in the following figure:

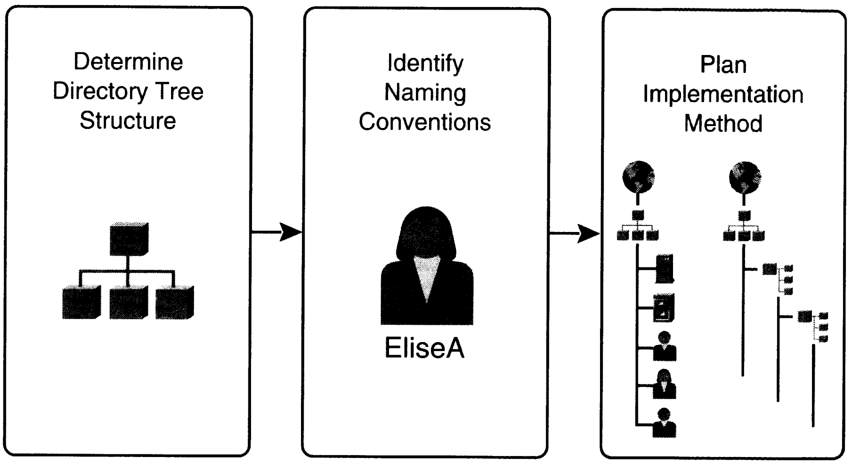
Determine Directory Tree Structure
Identify Naming Conventions
EliseA
Plan Implementation Method

Figure 2-2: Structural Design

A simple design does not require all the steps in both the structural and detailed design. A small installation using defaults can be performed without having to understand all the critical factors.

Detailed Design

The detailed design focuses on how the Directory is accessed by users, stored on the servers, and coordinated to provide accurate information. The steps in the detailed design are shown in the following figure:

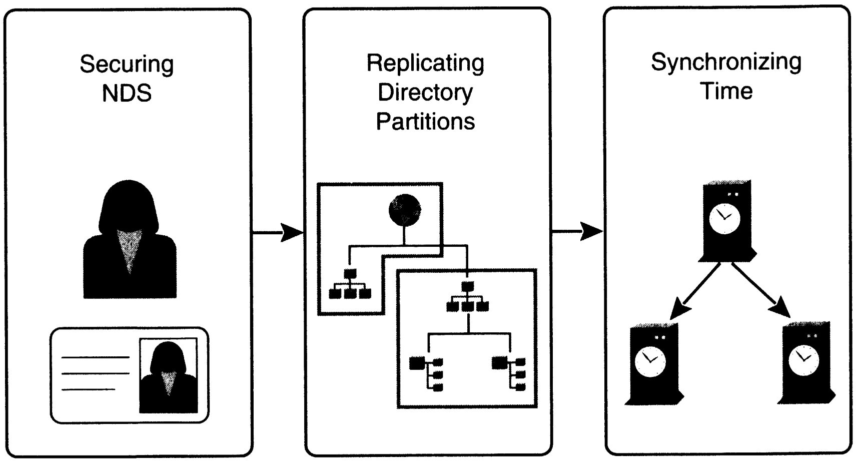

Figure 2-3: Detailed Design

This section focuses on the structural design of your Directory tree. You will create a detailed design after you learn about security, partitions, replicas, and time synchronization in the following sections:

- Section 3, "Securing the Directory Tree"
- Section 4, "Partitioning and Replicating NDS and Synchronizing Time"

Section 5, "Creating a Detailed Design and Troubleshooting NDS," contains an exercise in which you create a detailed design that includes information learned from Sections 3 and 4.

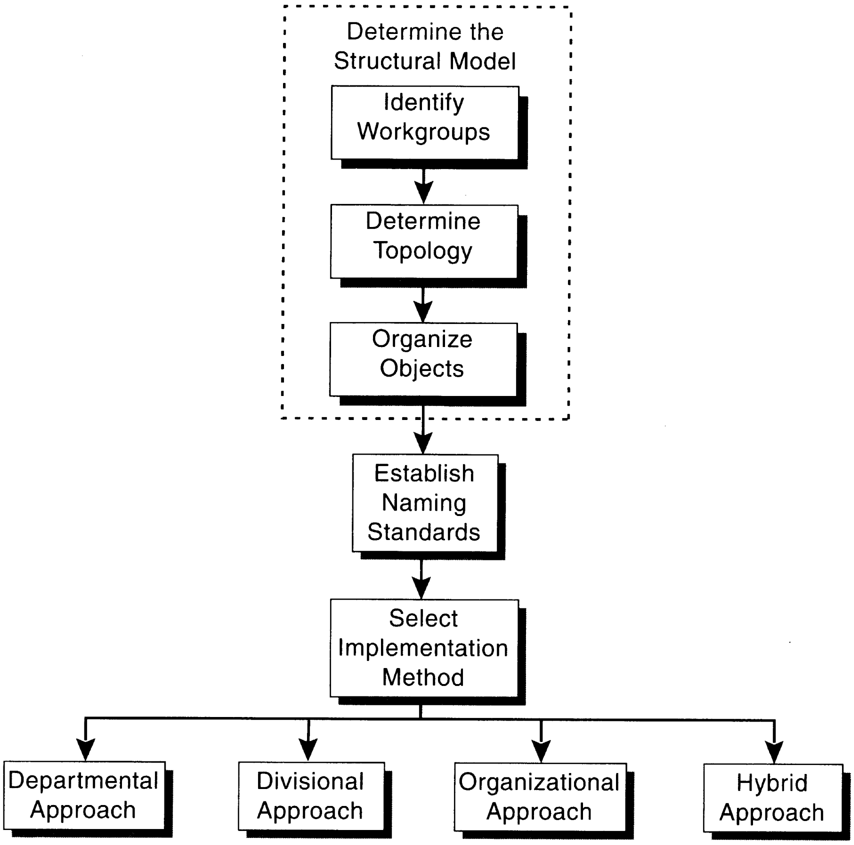

Figure 2-4

Creating the Structural Design

You should perform the following tasks to determine the structural design for your Directory tree:

1. Determine the structural model.
2. Establish the naming conventions.
3. Plan the implementation method.

Design Step 1: Determine the Structural Model

The model you base your design on can be created from several sources. These sources include network topology, workgroup needs, information flows, and resource access. Use organization charts, project team descriptions, and workflow analysis information to plan your model.

Identify Workgroups

Identifying workgroups is critical in designing an effective Directory tree. The Directory tree should be designed according to workgroup needs and the management of these workgroups. You should consider two important questions as you identify workgroups:

- How do users perceive themselves and their resources in their work environment?
- How are their resources accessed?

An effective Directory tree supports users in completing their tasks. It should uniquely identify all resources in your network as users see them.

Because users identify themselves with workgroups, organizing resources by workgroup is an effective way of designing the Directory tree. You can base your organization on one or more of the following groupings:

Administrative Divisions

In this model, the tree resembles the company's organization chart and emphasizes workgroups based on management structure.

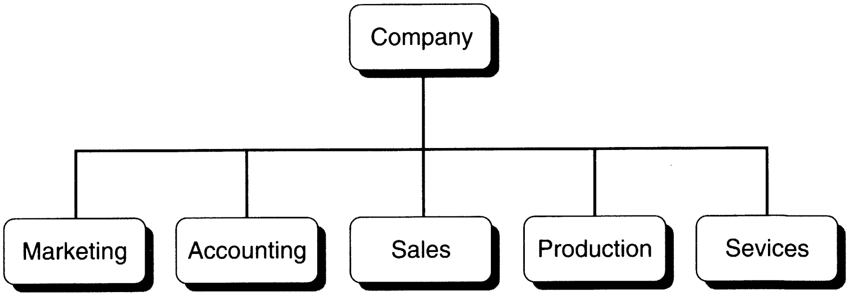

Figure 2-5: Administrative Model

Workgroups across Divisions

In this model, the tree combines elements of the administrative approach and the workgroup approach. It emphasizes workgroups based on similarities in users' tasks and resources.

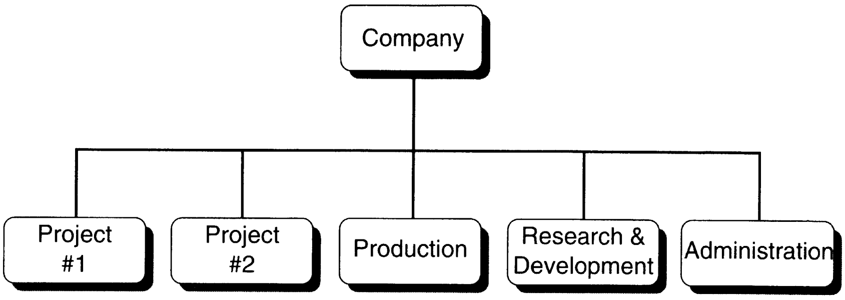

Figure 2-6: Workgroup Model

Geographical Locations

In this model, the tree resembles the company's physical layout and emphasizes workgroups based on physical location.

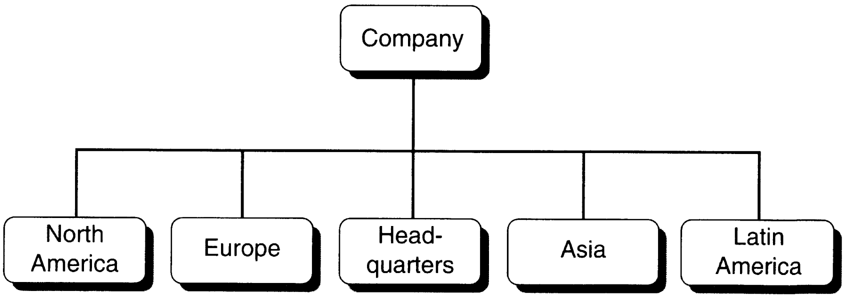

Figure 2-7: Geographical Model

Hybrid or Mixed Environment

Using a hybrid model, your tree structure could be a combination of the three models.

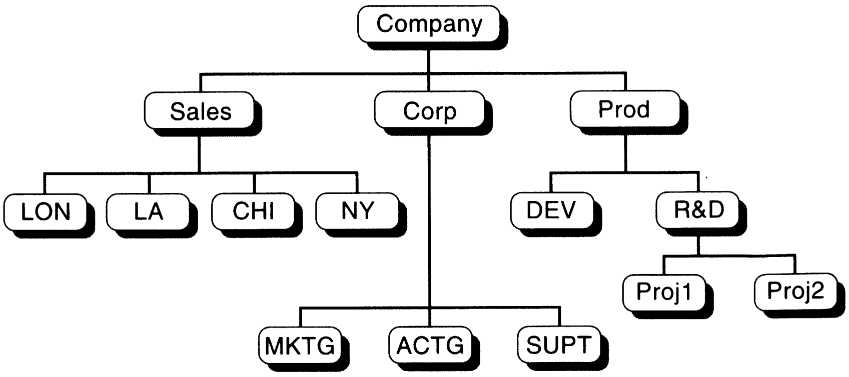

Figure 2-8: Hybrid Model

It is important to remember the main purpose of NDS: to model the network after a common human paradigm, making the network easy to use. Therefore, in choosing an approach, you should select the approach that will make sense to users and administrators of the network.

Once you acquire the information from the organization charts, network topology, and workgroup needs, you can start organizing the network resources using the NDS objects.

Determine Network Topology

A helpful resource in determining the model is your current network setup. Where are the WAN links in your network? How are network resources divided? You can use this information to help determine the structure of your Directory tree.

Organize Objects

Determine the number of layers when using the container objects to provide structure to the Directory tree. The Directory tree works well with approximately five to eight layers.

The Directory can handle any number of levels; however, the limitation comes when entering the path of an object in the name. The name with the path cannot be longer than 256 characters. The details of NDS naming will be presented later in this section.

Design Step 2: Establish Naming Standards

Novell® recommends that, before you implement your Directory tree, you create a document that describes your naming standards. The document should include information such as the following:

- The conventions you use for naming Directory objects, including users, printers, print queues, and servers
- The property values (such as telephone numbers and addresses) you use for the objects

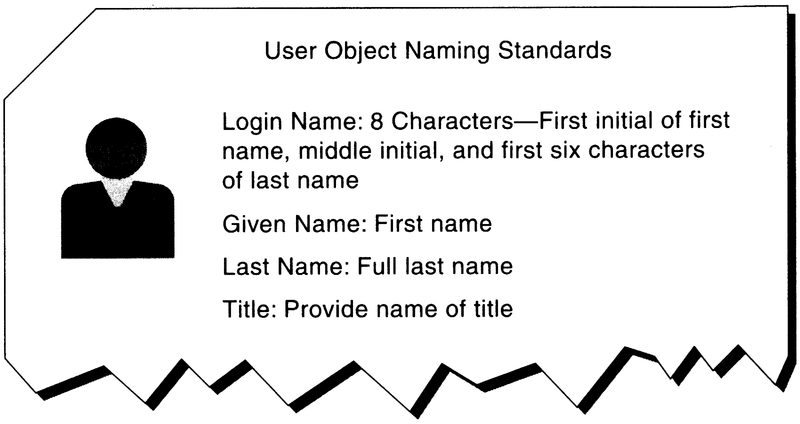

User Object Naming Standards

Login Name: 8 Characters—First initial of first name, middle initial, and first six characters of last name

Given Name: First name

Last Name: Full last name

Title: Provide name of title

Figure 2-9: Naming Standards

Understanding the Importance of Consistency

Consistency, especially in the naming schemes used for objects, provides several benefits:

- Consistent naming schemes provide a guideline for network administrators who will add, modify, or move objects within the Directory tree.
- Having the naming standards eliminates redundant planning. The standards give administrators an efficient model to meet their needs, but leave implementation of resource objects open and flexible.
- Consistent naming schemes help users identify resources quickly, maximizing users' productivity.

For further information about naming standards, see Appendix C in *Introduction to NetWare Directory Services.*

A practical example of applying naming schemes is usernames. A common guideline is to assign usernames using the first initial of the user's first name and up to seven characters of the user's last name. For example, user Thomas Jones would have a login name of TJones. Thomas can easily remember who to log in as and co-workers can easily predict other usernames. If another person named Tammy Jones joined the department, you could use a middle initial to make the username unique.

Devising Naming Conventions

The most important recommendation for Directory object names is that they be short but descriptive. This makes object names easy to remember and identify. This also makes Directory searches easier for users who do not know all the specifics of a resource they need to use.

For example, if Thomas wants to print to an Apple LaserWriter and he knows his organization represents this type of printer by starting the name with APL, he can narrow his search quickly.

Design Step 3: Select the Implementation Method

Once you have your Directory tree structure and a plan to implement consistent naming, you must determine your implementation method. To start this process, ask yourself questions similar to the following:

- Do you need to create a small Directory tree for a workgroup? If so, select the departmental method.
- Should you create a few large Directory trees for the company? If so, select the divisional method.
- Can a Directory tree be created at once for the entire company? If so, select the organizational method.

Departmental Implementation Approach

A departmental implementation creates a small Directory tree organized around a workgroup.

A departmental implementation is generally necessary in the following circumstances:

- Coordination of administrators, planners, and installers is not practical.
- No communication links are available to allow servers to access the same Directory over a WAN.

The departmental approach allows a group or department within the organization to implement NetWare 4 without waiting for the organizational networking goals to be determined by a central committee or IS department.

You must ensure unique Directory tree names if you have several trees on your network.

This approach allows smaller groups within an organization to experience immediate benefits from NDS by installing multiple Directory trees. NetWare 4 can accommodate these workgroups while maintaining the option to combine Directory trees later.

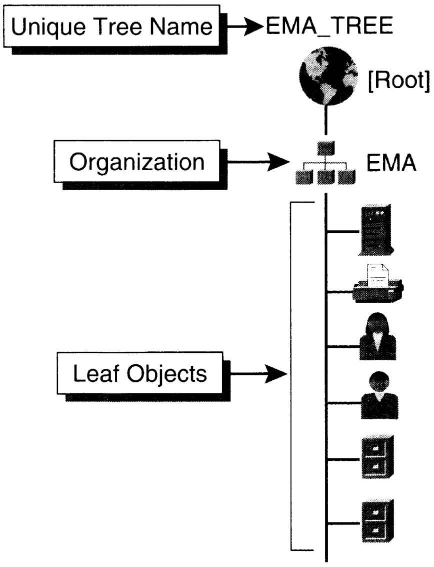

Figure 2-10: Departmental Directory Tree

If you plan to merge a departmental tree into a larger tree later, you should modify the structure slightly. Figure 2-11 shows how to change the structure to facilitate merging.

Adding the extra Organizational Unit allows you to move the container easily in the newly merged tree.

For each tree that will be merged, you must have unique Organization container names.

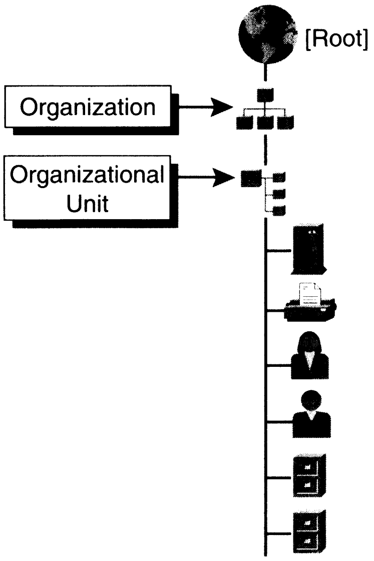

Figure 2-11: Departmental Directory Tree for Merging

Divisional Implementation Approach

The divisional approach creates a Directory tree for numerous workgroups, but not for the whole company. The divisional approach could use the location model (regional), the administrative model (a division), or the functional model (an entire subsidiary). For a large company, you could have a tree in each location. You could also merge the trees later, if needs demand an organizational implementation.

Figure 2-12 below shows a divisional Directory tree for the case company, EMA. Note that other divisional Directory trees may be created for the other divisions of EMA.

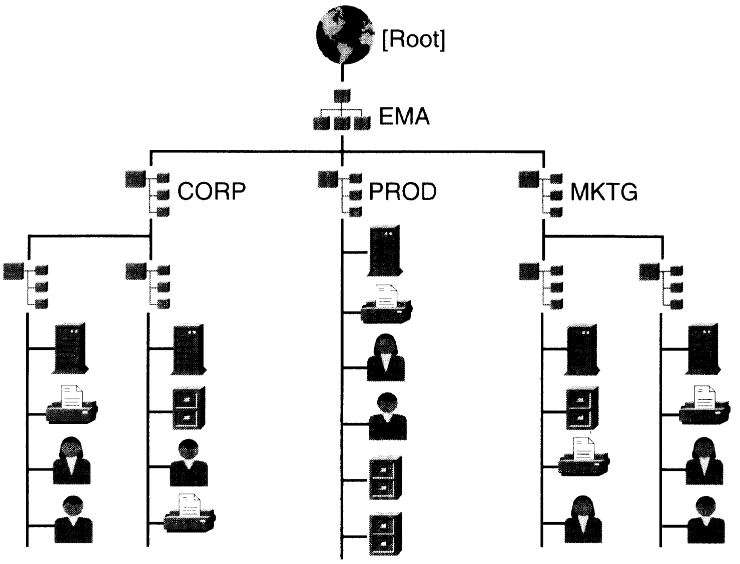

Figure 2-12: Divisional Directory Tree

Organizational Implementation Approach

Often referred to as a top-down approach, the organizational approach requires that you have a clear understanding of your organizational networking goals.

An organizational implementation creates one Directory tree. This type of implementation is generally possible when all servers are connected to each other, via either a LAN or a WAN. Other factors which suggest that an organizational implementation would work are as follows:

- The network is small enough and simple enough that you can upgrade to NetWare 4 without impacting the company's business.
- A centrally controlled group of administrators, such as an IS department, can manage the upgrade to NetWare 4.

Often, a centralized committee decides on the structure of the Directory tree's upper layers. Individual administrators design the tree's lower branches.

Workgroups identified by organization might look similar to the figure below. A hybrid structure with geographic locations at the top of the tree is a common implementation.

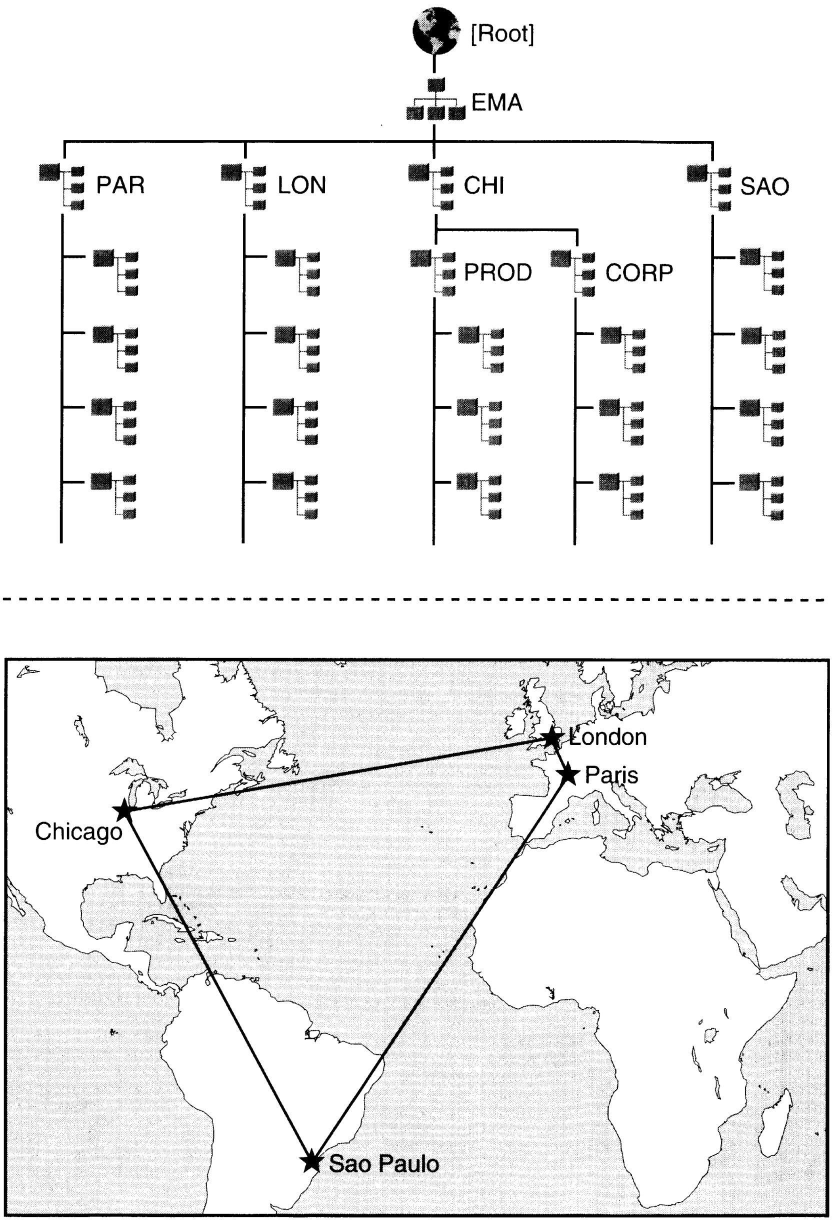

Figure 2-13: Organizational Implementation Approach

Review Structural Design Guidelines

The following table reviews the structural design guidelines:

[illegible]	[illegible]
1. Determine the Structural Model	❑ Identify workgroups: ■ Administrative divisions ■ Workgroups across divisions ■ Geographical locations ■ Hybrid models ❑ Organize objects.
2. Establish Naming Standards	❑ Importance of consistency ❑ Naming with bindery services ❑ Devising naming conventions
3. Select the Implementation Method	❑ Select an implementation: ■ Departmental ■ Divisional ■ Organizational

Table 2-1: Structural Design Guidelines

Directory Tree Examples

The next two pages present examples of Directory trees. The examples are categorized by structural model and implementation method. Each Directory tree focuses on the structure; consequently, you will see mostly container objects.

The hybrid Directory tree structure is not shown. It would be a combination of the models in the divisional and organizational examples.

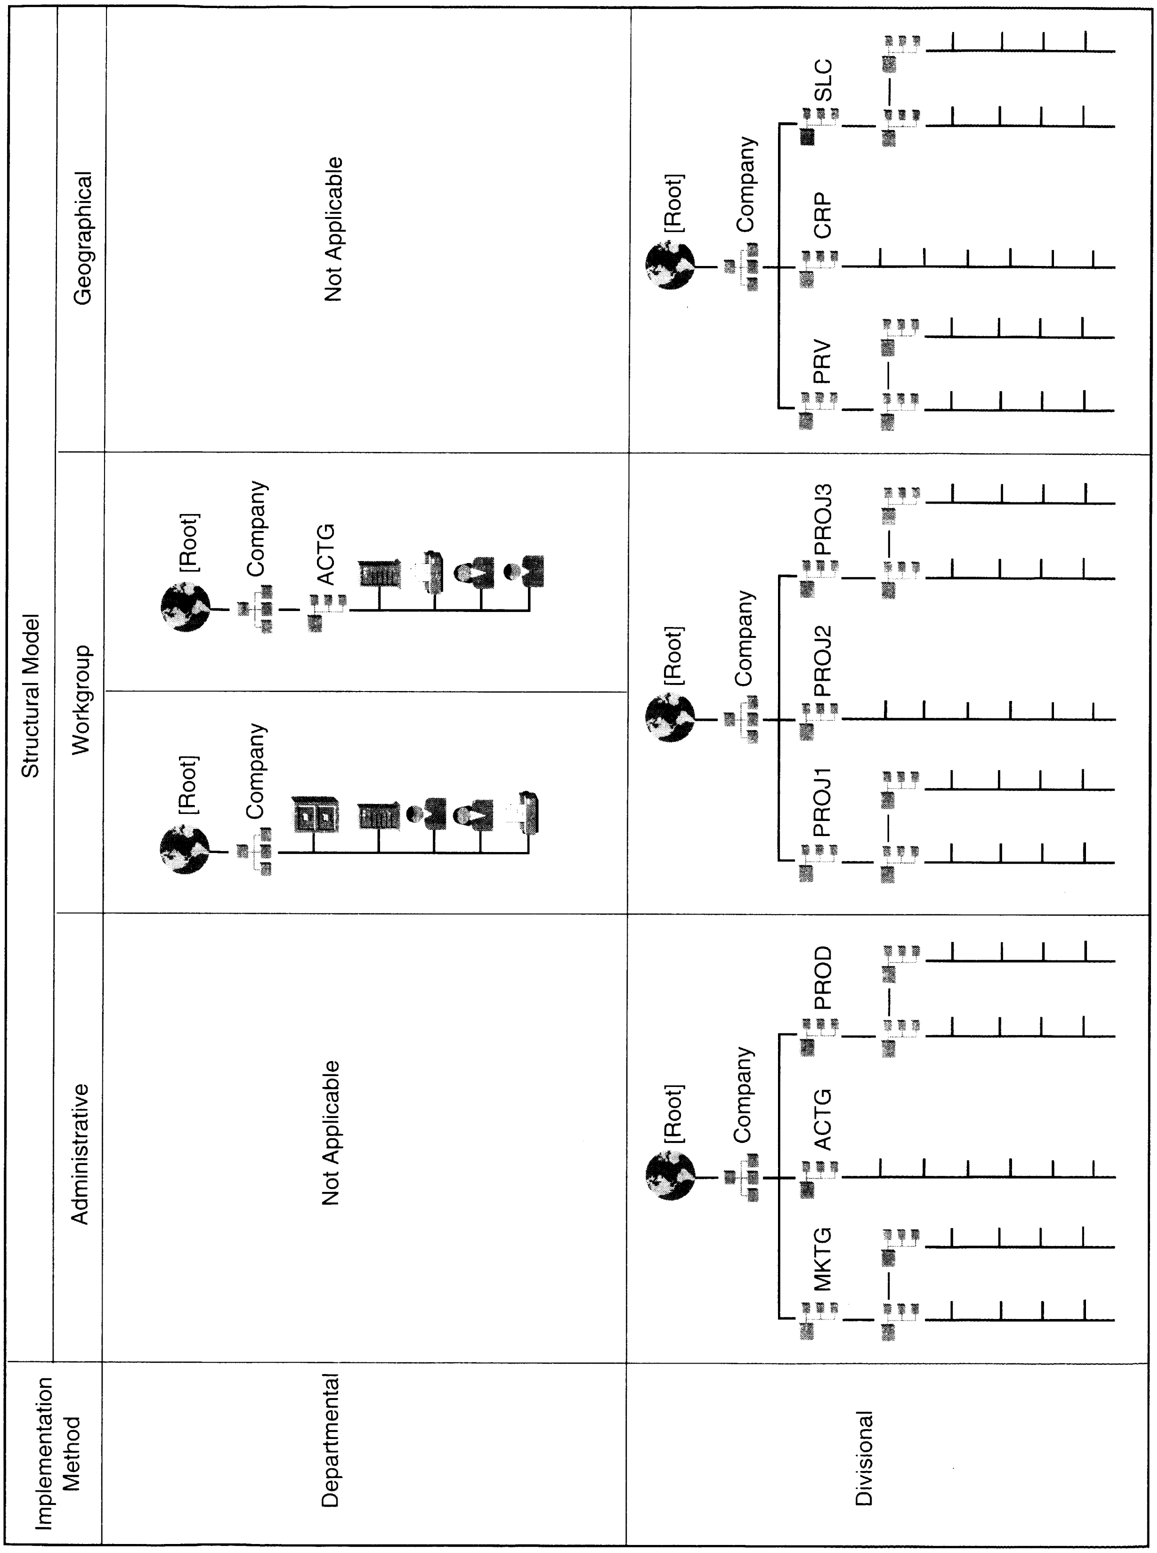

Figure 2-14: Directory Tree Examples

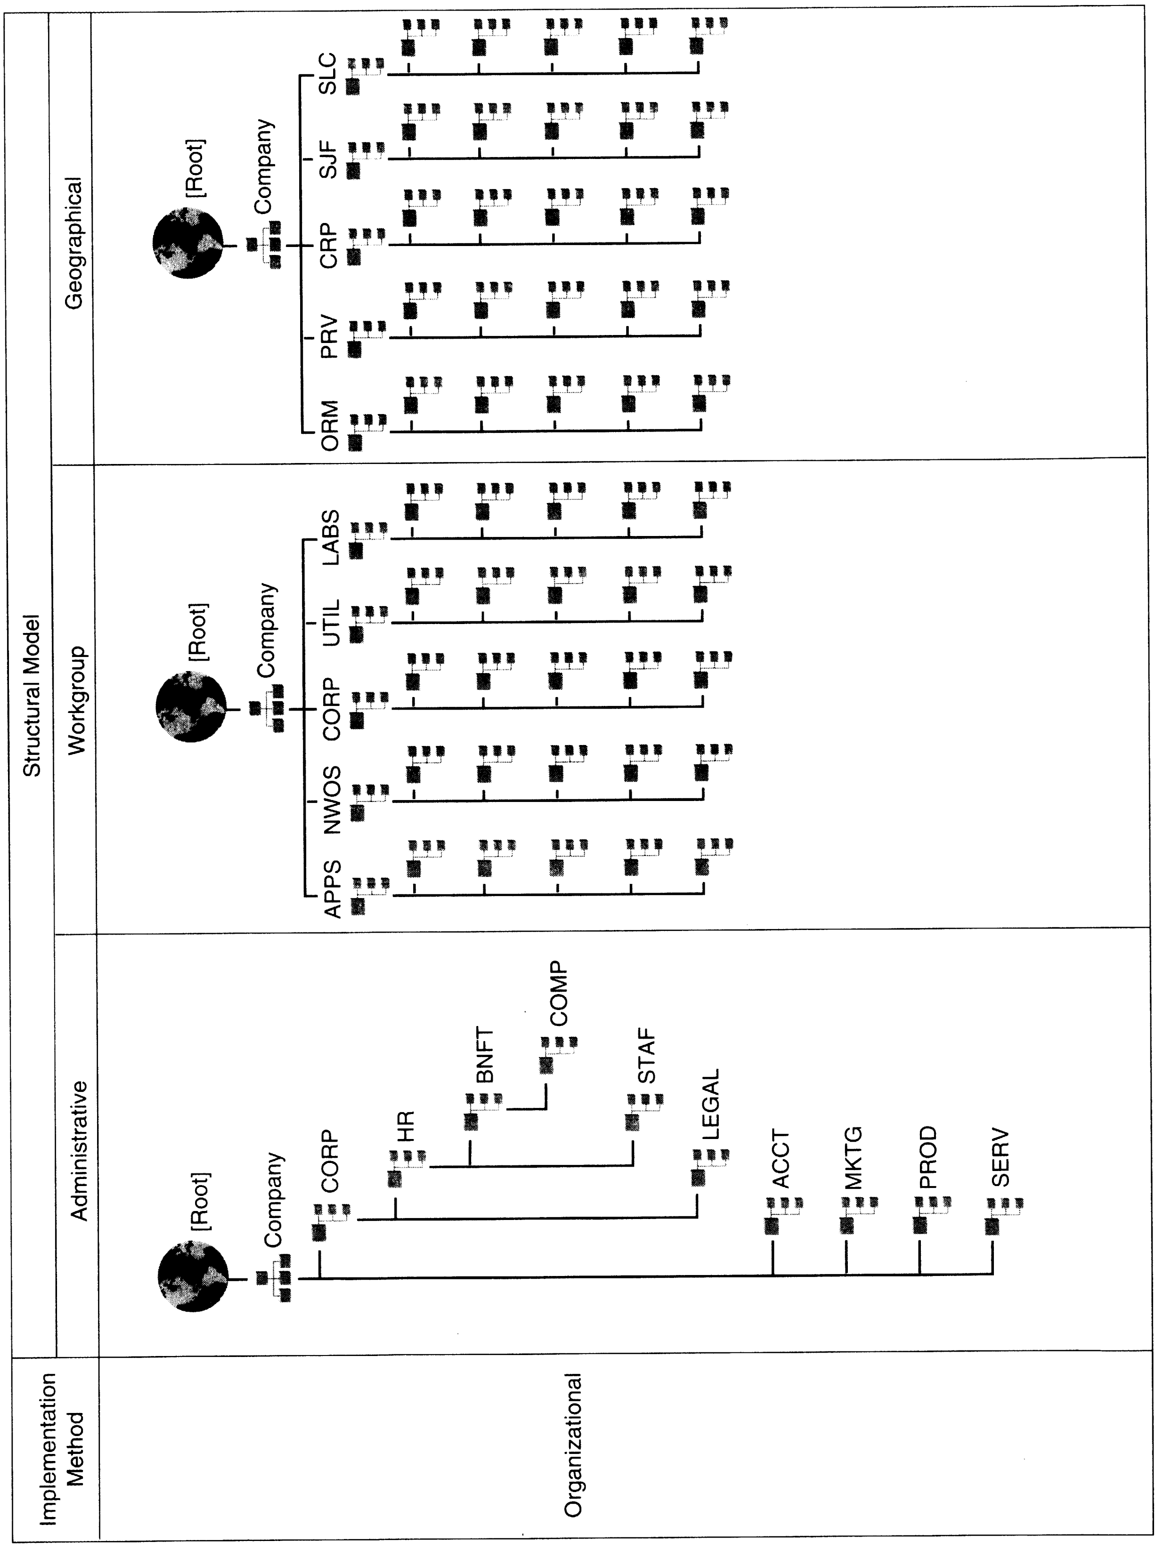

Figure 2-15: More Directory Tree Examples

Written Exercise 2-2: Design a Directory Tree

Activity: Design a Directory tree from the information provided.

Procedure: Using the assigned scenario, design a Directory tree and implementation to represent the organization's network resources and users.

1. The students will be assigned to one of three groups.
2. Each group will create a design and an implementation plan based on the assigned scenario.
3. Each group will identify the model and implementation approach and then design the Directory tree using the information assigned. Use the examples on the previous pages.
4. Each group will present the information and the associated rationale to the class.

When you have completed the exercise, save your Directory tree design for use in Exercise 5-1.

Scenario 1: TCabinets has headquarters in Toronto, Canada, with other sites in

- Milwaukee, Wisconsin (USA)
- London, England
- Bonn, Germany

Production is in Milwaukee and London.

Finance, Marketing, and Administration are in Toronto.

The International Marketing office is in Bonn.

Each location has a sales staff.

Currently, each location maintains its own server and has a communication link back to Toronto. The Finance department has two servers. Departments in each location have their own printers.

Scenario 2: Standard Publishing is a worldwide company. The MIS department has decided to implement individual Directory trees at each site.

You are the administrator for corporate headquarters in New York City. The departments at your location include

- Marketing
- Accounting
- Production
- Central Administration
- Development

Each group has at least one server and several printers.

Scenario 3: Sports, Inc. is an athletic equipment company located in Los Angeles, California. You are the network administrator for the company. After careful planning, you have determined that you will upgrade your current NetWare 3 servers to NetWare 4. You will create independent trees for each workgroup and then merge them into an organization-wide Directory tree.

Your workgroups include management, production, and sales. Each group has one NetWare 3 server. What will each Directory tree look like before and after the merge?

Merging Directory Trees

NetWare provides the ability to merge trees. If you start your design with the departmental approach, you can change the implementation to a divisional or organizational approach by merging the trees.

The case company, EMA, has decided to use this implementation method. The network administrators decided to use the geographical model to structure the network resources.

Each site has performed an upgrade from NetWare 3 to NetWare 4. The sites will merge into the corporate Directory tree, EMA_TREE. Figure 2-16 shows the current tree structure for each location.

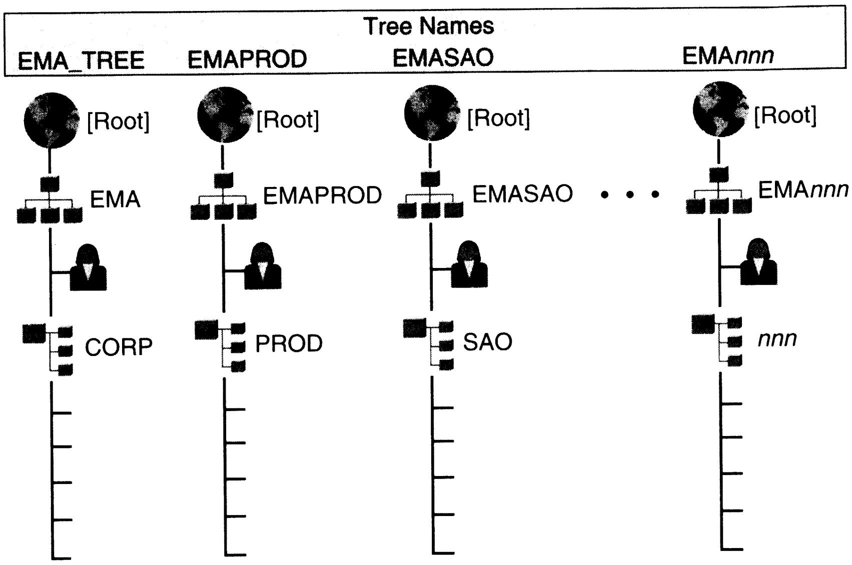

Figure 2-16: EMA Directory Trees

Tools Required

To merge the two trees, network administrators need the following tools (see Figure 2-17):

- The SBACKUP utility to back up the Directory
- SET TIMESYNC commands to change time synchronization parameters
- The DSMERGE utility to consolidate the trees
- The NetWare Administrator utility to customize the consolidated tree

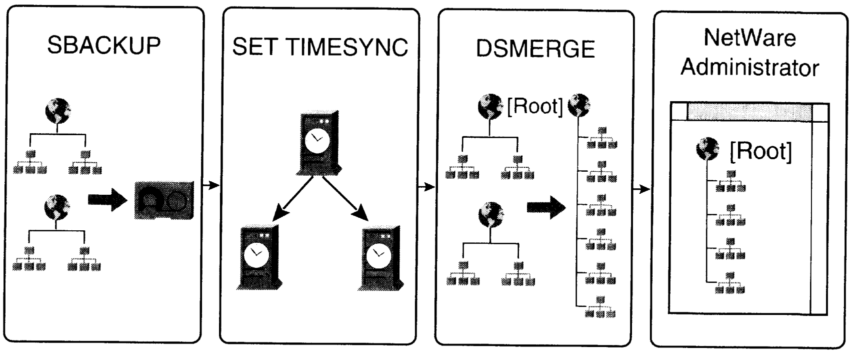

Figure 2-17: Tools for Merging Trees

Steps Required

These are the steps for merging two Directory trees:

1. Plan how to consolidate the trees.
2. Back up the Directory.
3. Establish time synchronization.
4. Run DSMERGE to consolidate the trees.
5. Clean up the consolidated tree.

Step 1: Plan How to Consolidate the Trees

You need to plan how you will consolidate the trees. Before merging the trees, you must answer these questions:

- Which tree will you consolidate into the other?

  Select the source tree and target tree. The source tree should be the tree with fewer objects at the root.

- What will the new tree look like?

  The two departments must agree on what the consolidated tree will look like, even though NDS flexibility allows them to continue to work independently.

Figure 2-18 below shows the source and target trees.

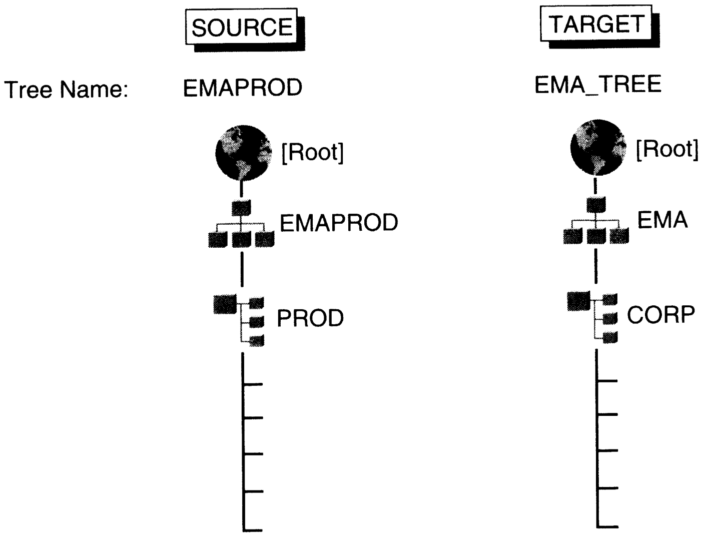

Figure 2-18: Source and Target Directory Trees

A merge forces a major structural change in the Directory. To ensure a smooth merge, perform the following tasks:

- Choose the source and target Directory trees.
- Create unique Organization object names for the source tree.
- Gather the full name for the Admin object and password for both the source and target trees.

Step 2: Back Up the Directory Tree

You should back up the Directory tree. This will ensure that you can start the merge again in case you encounter problems during the process.

Step 3: Establish Time Synchronization

The two trees need to agree on the correct time in order for DSMERGE to run. Time synchronization will be discussed in detail in Section 4, “Partitioning and Replicating NDS and Synchronizing Time.”

NDS depends on time to make updates to the Directory. Each server depends on a time source to track transactions, such as noting the time when a file is created.

As you coordinate time across several servers in the Directory, keep in mind the following rules:

- Not all servers can provide time; some must be time providers and time receivers.
- Time must be coordinated with all the servers before merging.

NetWare 4 has four different types of time servers. The two discussed here will enable you to merge the tree. The other types are covered in Section 4, “Partitioning and Replicating NDS and Synchronizing Time.”

Single Reference servers are the default type of time server. You cannot have more than one Single Reference server on a Directory tree because this type of server assumes it always has the correct time. All the other servers in the Directory tree must be time receivers.

Secondary time servers look to other servers for time information. Before you merge trees, you need to change all the servers, except the target tree server, to be secondary time servers. You also need to tell the Secondary servers to get time from the Single Reference server (see Figure 2-19).

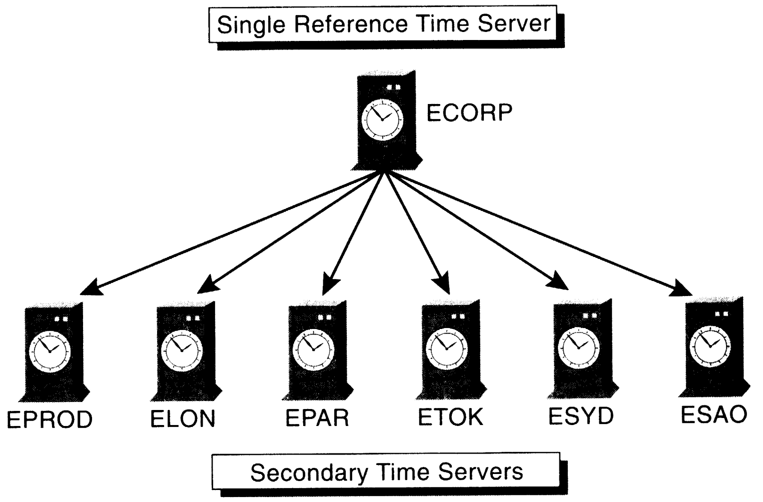

Figure 2-19: Coordinating Time Servers

To make each server a time receiver and ECORP the time source, type the following SET TIMESYNC commands at the server console:

1. **SET TIMESYNC TIME SOURCE = ECORP** <Enter>
2. **SET TIMESYNC TYPE = SECONDARY** <Enter>
3. **SET TIMESYNC WRITE PARAMETERS = ON** <Enter>

   This command creates the TIMESYNC.CFG file with the new parameters.

Type LOAD EDIT TIMESYNC.CFG to make sure that it includes the following: **Type = SECONDARY** and **Time Source = ECORP.**

4. **SET TIMESYNC RESTART FLAG = ON** <Enter>

   The RESTART FLAG command causes the server to read the changes from the commands in the configuration file immediately.

When the commands are successful, you will receive the following message:

```
Time synchronization has been established.
```

You can now merge both trees.

Step 4: Run DSMERGE to Consolidate the Trees

After you have backed up the Directory tree and coordinated time synchronization, you are ready to merge the trees.

The DIrectories for both trees must use the same schema. (The schema defines object classes and rules of containment.) If one of the Directories uses a modified schema, or if different versions of NDS are represented, you must correct this difference before merging trees.

To merge the trees, you need to load DSMERGE at the server console and then perform the following steps:

1. Check the servers by selecting the **Check servers in this tree** option.
2. Check the time synchronization by selecting the **Check time synchronization** option.
3. Merge the trees by selecting the **Merge two trees** option.

   You can use the DSTRACE screen on the target server to monitor the process. Make this screen active by entering the following command at the server console: SET DSTRACE = ON. Toggle to this screen by pressing <Alt>+<Esc>.

   When you select the **Merge two trees** option, you will see a screen that describes the merge process. After checking the two trees, a prompt asks if you want to merge. Then a screen shows the process of the trees merging. After the process, you will see a final screen that tells you the merge is complete.

If you are merging more than two Directory trees, make sure you allow enough time for each merge to finish before starting the next merge. Merging too soon can corrupt the target Directory tree.

4. Confirm the merge by launching NetWare Administrator on a workstation. If you already have NetWare Administrator running, you need to collapse and expand the [Root] to see the new organization in the tree.

Step 5: Clean Up the Consolidated Tree

If the objective is to merge the trees into one organization, you may need to clean up the consolidated Directory tree. Figure 2-20 shows the Directory tree immediately after the merge.

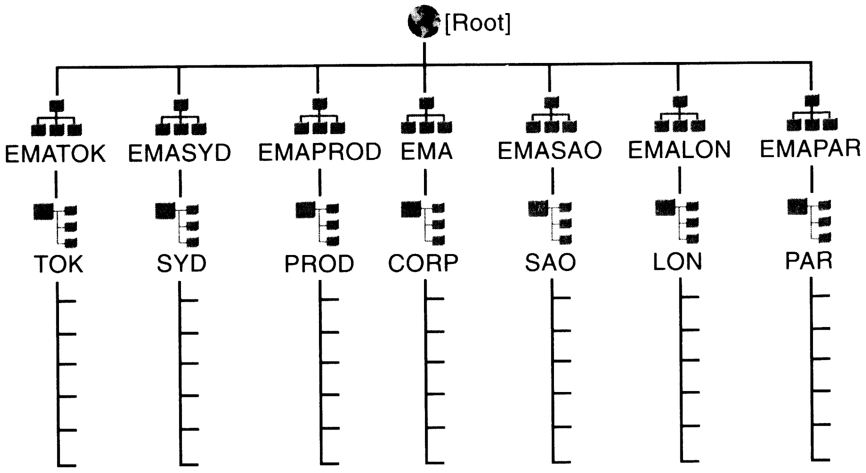

Figure 2-20: Merged Directory Tree

To begin cleaning up the Directory tree after a merge, complete the following steps:

1. Move containers with their objects.

   The process of moving containers is covered in Section 4, "Partitioning and Replicating NDS and Synchronizing Time."

2. Rename objects.

   Rename the objects according to your naming conventions.

The consolidated EMA tree is shown in Figure 2-21.

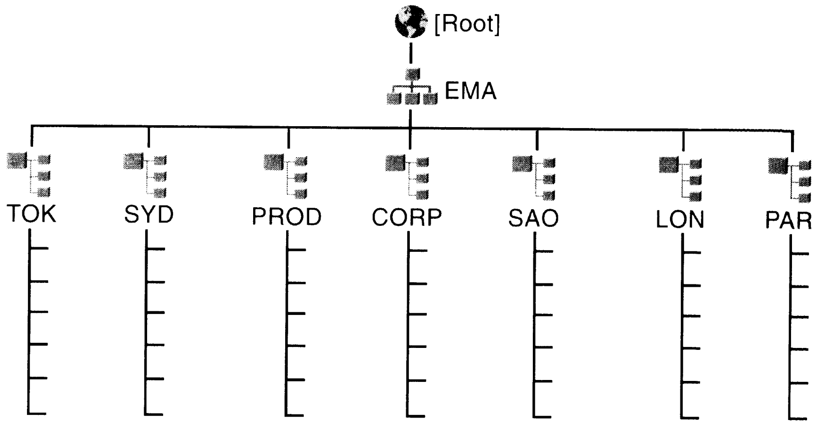

Figure 2-21: Directory Tree after Moving Containers

The following are optional tasks you can complete to clean up the consolidated Directory tree.

- **Ensure that login scripts exist for containers in the new tree.** Copy any container login scripts in the existing containers to the new containers. If you are using the NWADMIN utility, you can use the MS Windows cut (Ctrl-X) and paste (Ctrl-V) functions.
- **Delete unnecessary containers.** Delete the container after you have transferred all the objects and their required information.

Implications from the Merge

Once you have a combined tree, you need to change parameters on the server and the client. The names have now changed because of the new context. You need to change references to these names on the server and client. Examples of parameters that need to be changed include the following:

- NAME CONTEXT = (in NET.CFG on the workstations)
- PREFERRED TREE = (in NET.CFG on the workstations)
- BINDERY CONTEXT = (in AUTOEXEC.NCF on the servers; bindery context will be explained in Section 6.)

Hands-On Exercise 2-3: Merging Directory Trees

Activity: Merge the six departmental Directory trees into the EMA_TREE target Directory tree.

Procedure: These are the phases to merging the Directory trees:

- Plan and coordinate the merge of all trees involved.
- Back up the Directory trees.
- Change time synchronization configuration to allow the trees to be merged.
- Merge the Directory trees.
- Confirm that the merge is complete.

Plan the consolidated corporate tree. Coordinate the order of when trees will be merged.

1. Plan and draw the consolidated corporate tree. Make sure you merge at the [Root] level.

2. Meet with the other administrators to decide the order in which they will merge their trees. Write the order below.

   a.

   b.

   c.

   d.

   e.

   f.

Change the time synchronization type on your server.

This process coordinates time on the Directory trees prior to merging.

DSMERGE does not allow Directory trees to merge if time is not synchronized across the two Directory trees.

1. At the server console, type the following to add the time source:

   SET TIMESYNC TIME SOURCE= ECORP <Enter>

2. Type the following to change your time server:

   SET TIMESYNC TYPE = SECONDARY <Enter>

3. Type the following to create the TIMESYNC.CFG file:

   SET TIMESYNC WRITE PARAMETERS = On <Enter>

4. (Optional) Look at the TIMESYNC.CFG file to ensure the parameters are correct.

   a. Type

      LOAD EDIT TIMESYNC.CFG <Enter>

   b. Make sure the following parameters have been included:

      TYPE = SECONDARY

      TIME SOURCE = ECORP

   c. Press <Esc> twice to exit.

5. Type the following to make the server read the TIMESYNC.CFG file immediately:

   SET TIMESYNC Restart Flag = On <Enter>

   Within moments, you should get the following message:

   ```
   Time synchronization has been established.
   ```

   Now you are prepared to merge Directory trees.

Merge the Directory trees.

The Directory trees must be merged one at a time.

Perform the following steps to merge your Directory tree into the corporate Directory tree.

1. At the server console, type

   LOAD DSMERGE <Enter>

2. Select **Check Time Synchronization**.

   Verify that your server type is SECONDARY.

3. Press <Esc> to return to the main menu.
4. Select **Merge two trees**.
5. For "Administrator's name," type

   ADMIN.EMA*nnn* <Enter>

6. Enter your password.
7. Select the field next to the title "Target tree:" and press <Enter>; then select **EMA_TREE.**
8. Select the field next to the title "Administrator name" and type

   ADMIN.EMA <Enter>

9. Ask your instructor to type in the administrator's password for the EMA_TREE tree.
10. Press <F10> to start the merge.
11. Read the information screens and answer **Yes** to the confirmation prompts.
12. Exit DSMERGE

You will get a message that the merge is complete.

13. Use DSREPAIR at the target server console to verify the merge is complete before executing the next merge:
    a. Type **LOAD DSREPAIR** at the target server console.
    b. Select **Replica synchronization.**
    c. Verify **Total errors = 0** in the upper right portion of the screen.

Confirm that the Directory trees are merged.

1. View the merged tree in NetWare Administrator.
2. Double-click twice on the [Root] to view the new tree structure.

   Double-clicking twice collapses and expands the Directory tree.

You will not perform any cleanup tasks at this time.

Hands-On Exercise 2-4: Navigating the EMA Directory Tree

Activity: This exercise will familiarize you with browsing the Directory tree. You will use the NetWare Administrator and CX utilities.

The CX utility is similar to the DOS CD command. The CX command, like the CD command, allows you to move around the tree structure. You can use CX to view the Directory tree and move to different contexts. For example, your Directory tree looks like this:

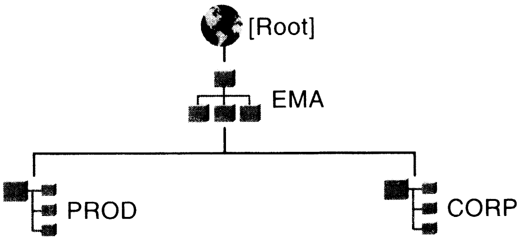

If your current context is OU=PROD, you would enter CX .OU=CORP.O=EMA to change to OU=CORP.

Part I

Procedure: Activate NetWare Administrator, perform the following tasks, and answer the following questions.

1. What container and its objects are displayed when you activate NetWare Administrator?

2. Select **Set Context...** from the View menu.

3. Enter **[Root]** in the New Context line.

4. How many objects are in the [Root]?

5. What object classes are represented directly under the [Root]?

6. Double-click on the object icon for the organization (EMA).

7. How many objects are in the Organization?

8. What object classes are represented directly under your Organizational Unit?

Part II

Procedure: Open a DOS window in MS Windows (or exit MS Windows) and perform the following tasks at the DOS prompt.

1. Change to the network drive.

It may be helpful to type **CX /T** to see containers, or **CX /T /A** to see all objects within a container, after each step.

2. Display your current context. Type

   CX <Enter>

3. Display the objects in your current context. Type

   CX /T /A <Enter>

4. Display the objects in the entire Directory tree. Type

   CX /T /A /R <Enter>

5. Create and display a file that contains the entire Directory tree structure. Type

   CX /T /A /R > TREE.TXT <Enter>

   TYPE TREE.TXT <Enter>

6. Change your current context to the [Root]. Type

   CX /R <Enter>

7. Display the container objects in the [Root]. Type

   CX /T <Enter>

8. Change context to the PROD Organizational Unit object. Type

   CX EMA <Enter>

   CX PROD <Enter>

9. View all objects in the PROD container with the CX command. Type

   CX /T /A <Enter>

10. Change context to the CORP Organizational Unit object one container at a time.

    a. To change context to the [Root], type

       CX /R <Enter>

       or

       CX .. <Enter>

    b. To change context to the CORP container, type

       CX EMA <Enter>

       CX CORP <Enter>

nicht unterstützt

11. Change context back and forth between PROD and CORP, using a single command each time.

    a. To move to PROD, type

       CX .PROD.EMA <Enter>

    b. To move to CORP, type

       CX .CORP.EMA <Enter>

Using Distinguished Names in Utilities

If your current context is the same as the context of the object you want to use, you can specify that object using its common name. If the object is outside your current context, however, you must provide NDS utilities with more information.

A *typeful distinguished name* exactly identifies an object and contains all the information NDS needs to locate the object in the Directory tree.

For example, consider the EMA Directory tree shown below. If you want to access the objects in .O=EMAPROD, any of the following commands would work, regardless of your current context:

```
LOGIN .CN=DOUGC.OU=PROD.O=EMA
CAPTURE P=.CN=EPROD_P.OU=PROD.O=EMA
NAME CONTEXT = "OU=CORP.O=EMA"
```

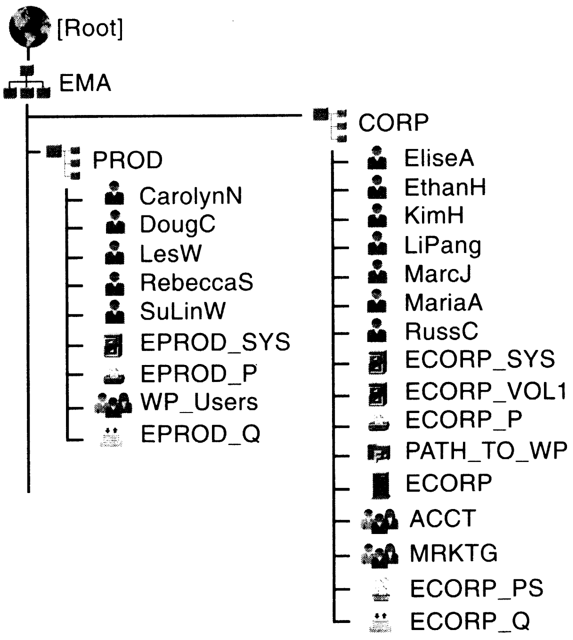

Figure 2-22: The EMA Directory Tree

Written Exercise 2-5: Supplying Object Names

Activity: Provide correct NDS names for a Directory tree structure.

Procedure: Use the Directory tree structure shown below to answer the following questions.

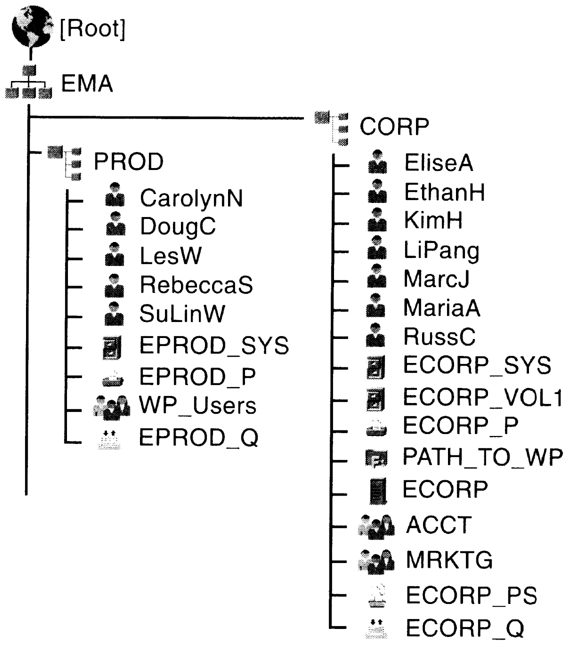

1. Provide a typeful distinguished name for User object LiPang.

   .CN=LiPang.OU=CORP.O=EMA

2. Provide a typeless distinguished name for Printer object EPROD_P.

   .EPROD_P.PROD.EMA

3. Write a typeful distinguished name for the Organizational Unit object CORP.

   O=
   .OU=CORP.EMA

4. If your current context is OU=CORP, what is the simplest name that accurately refers to the User object KimH?

   KIMH

5. Your current context is OU=CORP; write a relative distinguished name that accurately refers to the Volume object EPROD_SYS.

   .PROD
   EPROD_SYS..

6. Your current context is OU=PROD; write a typeless relative distinguished name that refers to the User object DougC.

   DougC

7. Describe current context.

8. Now that users in the PROD container have changed context, what typeless context would you enter with the NAME CONTEXT = command in the NET.CFG file?

   NAME CONTEXT = " OU=PROD.O=EMA "

If your user account is in the PROD container, make this change in your workstation's NET.CFG file so you can log in the correct context.

Creating User Objects with UIMPORT

NetWare 4 allows you to take advantage of user database files that are currently on your system. Rather than re-enter the user information manually when you set up NetWare 4, you can use the User IMPORT (UIMPORT) utility to add new User objects into the Directory tree, as shown in Figure 2-23. UIMPORT also gives you the capability to update User objects when common data is changed.

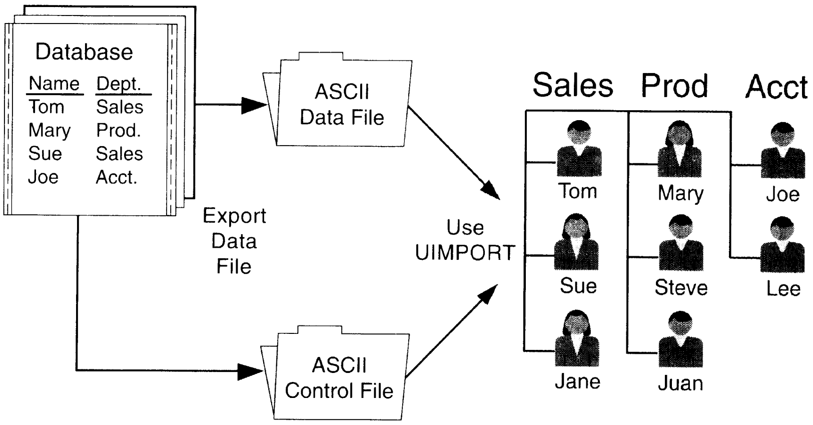

Figure 2-23: Creating Users with UIMPORT

To use UIMPORT, you must first export the database into a delimited ASCII file. You can export the database with most database programs and some spreadsheet programs. The default delimiter is a comma.

UIMPORT allows you to specify the format and field sequence for the data being imported. To do this, you must create an ASCII control file. The control file describes how the data is imported into the Directory.

You can set parameters (such as whether UIMPORT should create new objects or just update existing objects). You can also change the default delimiter to a different character.

You can set the context in the control file so that the users will be placed in a context other than the current context. To do this, you must use the complete context, without a leading period (.). If you do not specify a context in the control file, the users are placed in your current context.

For more information about UIMPORT, refer to *Supervising the Network.*

A data file may look like the sample below:

```
Jones,Adam,J,111 South 8th East,Salt Lake City,
Utah,84007,2345,Sophomore,Environmental
Engineering,2.8,Engineering Sciences

Smith,John,D,222 North Cerillos,Los
Angeles,California,96000,2875, Senior,Accounting,
3.0,Business Administration
```

A control file may look like the sample below:

```
Import control
    Name context=.administration.student_accts
    User template=y
    Create home directory=y
    Home directory path="Students/Home"
    Home directory volume=".SYS.Student Records"
Fields
    Last name
    Given name
    Middle initial
    Mailing label information
    Mailing label information
    Mailing label information
    Mailing label information
    Name
    Skip
    Skip
    Skip
    Department
```

Hands-On Exercise 2-6: Importing Users with UIMPORT

Scenario: You have a large group of users who need to be added to the Directory tree. You have created a text file with information about the users, and you are now ready to import this information into the Directory.

Procedure: The ASCII file containing information about the users is in the SYS:SETUP directory on your E*nnn* server. Refer to the instructions on the previous pages or to *Supervising the Network* for more information.

Perform the following:

1. From the DOS prompt, view the contents of the SYS:\SETUP\LIST.DAT file. Type

   MAP K:=E*nnn*_SYS:\SETUP <Enter>
   K: <Enter>
   TYPE LIST.DAT <Enter>

2. Write down the fields that you would use for the User object properties displayed in LIST.DAT.

Field

3. Using a DOS text editor, create a control file named LIST.CTL.

   You can open the UISAMPLE.CTL file in the SETUP directory and use it as a template to work from. You will need to delete some of the extra fields.

4. Save the file you created as LIST.CTL and exit the text editor.

5. From the DOS prompt, run the UIMPORT program with the control file and the data file as parameters. Type

   UIMPORT LIST.CTL LIST.DAT <Enter>

6. Launch MS Windows and NetWare Administrator.

7. Select the container to view the results of the import operation.

Summary

This section covered basic Directory tree design concepts. By designing the Directory tree before creating it, users can access the Directory more efficiently. NDS utilities allow you to change your Directory tree after it has been created.

You also learned how to implement tree designs for various approaches. For example, you can set up a Directory tree for a small department and then merge the tree into a larger organization.

Notes

SECTION 3 Securing the Directory Tree

Introduction

In this section, you will discover the similarities and differences between NDS security and file system security. You will identify the need for NDS security. You will also identify the default object and property rights and design a security strategy for a Directory tree and make trustee assignments.

Objectives

At the end of this section, you will be able to do the following:

1. Identify and explain the need for NDS security.
2. Perform steps to assign additional object and property rights.
3. Identify additional NDS rights users may require when using specific NDS objects: Alias of a User object, Profile Login Script, and Directory Map.
4. Given a Directory tree, determine effective rights for objects and properties and troubleshoot an NDS security scenario.
5. Set up an auditor User object and enable auditing.

For an overview of NDS security, see "Security" ("Trustees") in *Concepts*.

Exercise 3-1: NDS Security Review

This will be a group exercise. Your instructor will provide directions.

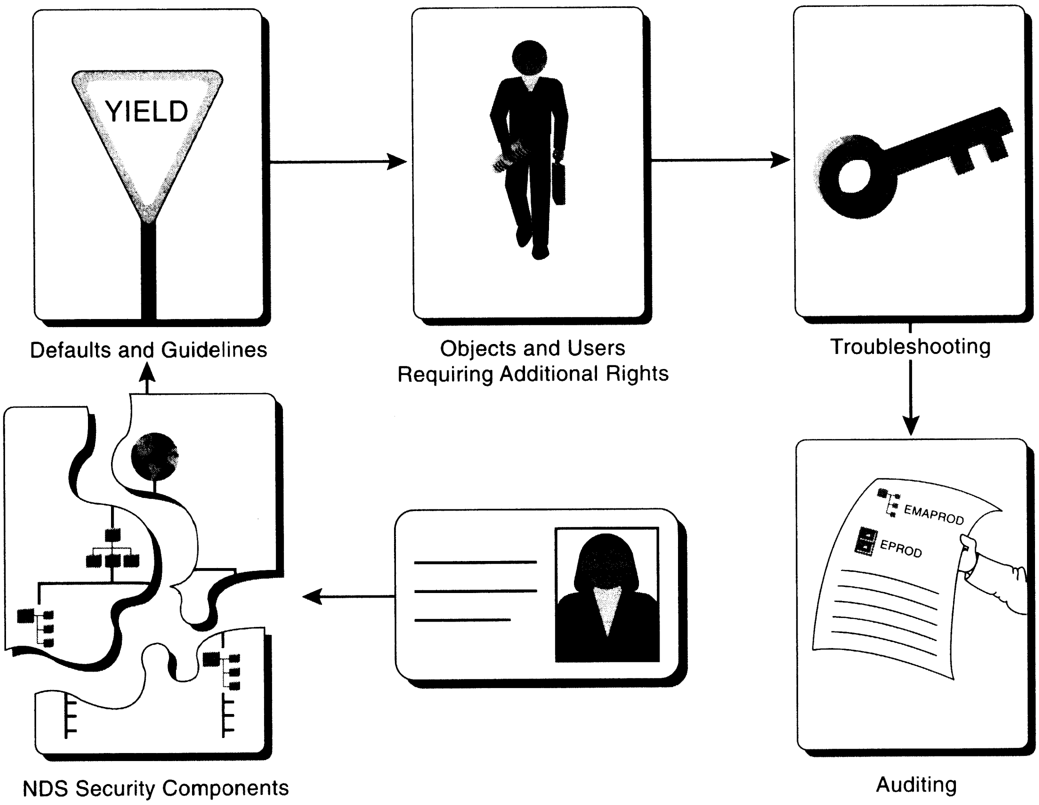

Figure 3-1

NDS Security Overview

As shown in the figure above, NDS security is separated into five topics which are presented in this course and in Course 520, NetWare 4 Administration. You should be familiar with the following topics from the NetWare 4 Administration course:

- NDS Security Components — concepts such as Object trustees, Object and Property rights, and inheritance.
- Defaults and Guidelines — default rights and guidelines for assigning additional rights.

In this course, you will learn about the remaining NDS security concepts and perform additional exercises for the following topics:

- Object and Users Requiring Additional Rights.
- Troubleshooting — This section identifies ways in which rights can be obtained and allows you to test your understanding of NDS security.
- Auditing — This section provides the basics of setting up an auditor who can monitor and report network activities.

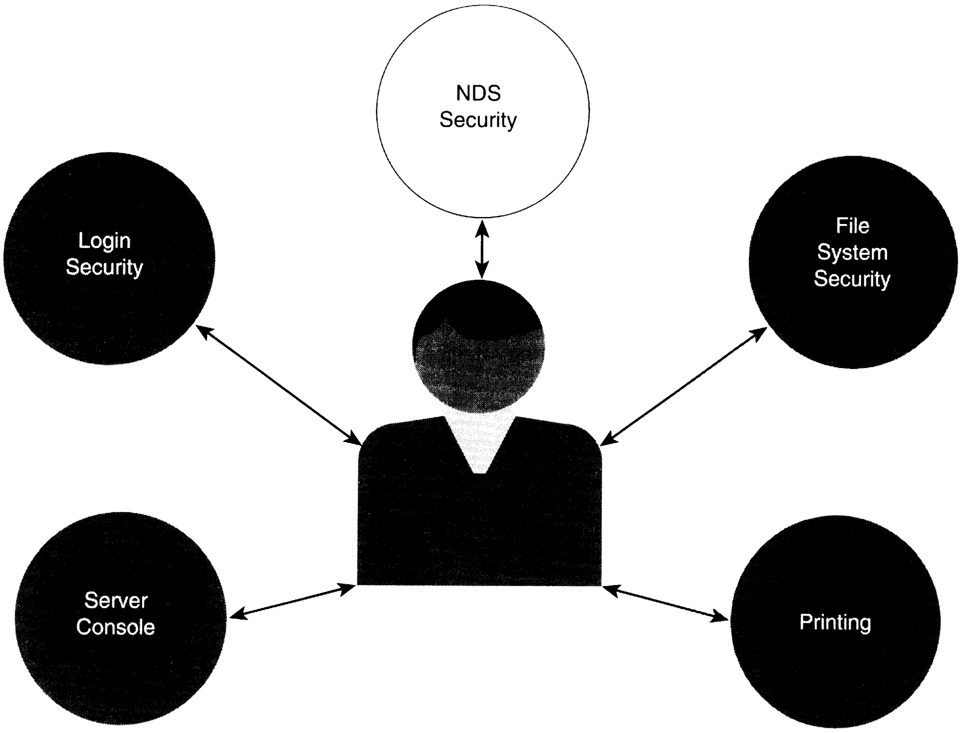

Figure 3-2

The Need for NDS Security

NDS expands the concept of a network. The sole definition of a network is no longer a single server, connected users, and resources. Now the network can be one server or fifty with resources that can be distributed throughout a building or throughout the world.

Because of the expansion of the network concept and the administration involved in such a network, NDS and NDS security are required to access and protect network resources and information about those resources.

Network Administration

One of the advantages of NDS is the ability to section off certain areas of the network and delegate network administration tasks, as in the following examples:

- The network administrator delegates administration tasks to container administrators or assistant administrators, providing more efficient administration of the network.
- In a government organization or another organization with high security requirements, the network administrator may only be involved in setting up the network. The network administrator then passes the administration rights to someone within the high-security workgroup. The network administrator's rights can then be blocked to protect sensitive information. The entire administration of that particular workgroup then becomes the responsibility of the container administrator.

Determining Additional Security Needs

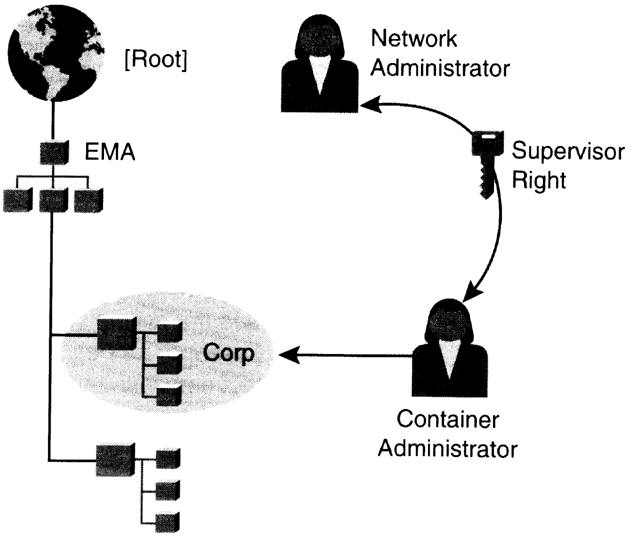

Figure 3-3: Additional Security Needs

As you can see, many of the default rights assignments are handled by NetWare 4. With NetWare 4 default NDS rights assignments, users created in the same container as the server can log in and use other system resources, such as print queues, defined in the same container. The default installation is very typical for organizations in which the users will be in the same container as their most frequently accessed servers.

Sometimes, however, you may need to assign additional NDS rights and file system rights. The following are typical examples:

- Dividing the tree's administration
- Assigning profile login scripts
- Creating Directory Map objects
- Allowing a user to log in with an alias in containers other than the container in which his or her original User object is stored
- Establishing special types of users such as mailing list administrators
- Assigning rights to facilitate the traveling user
- Determining NDS rights needed for printing
- Assigning file system rights to SYS:PUBLIC for a user whose User object does not reside in the same container as the desired SYS Volume object

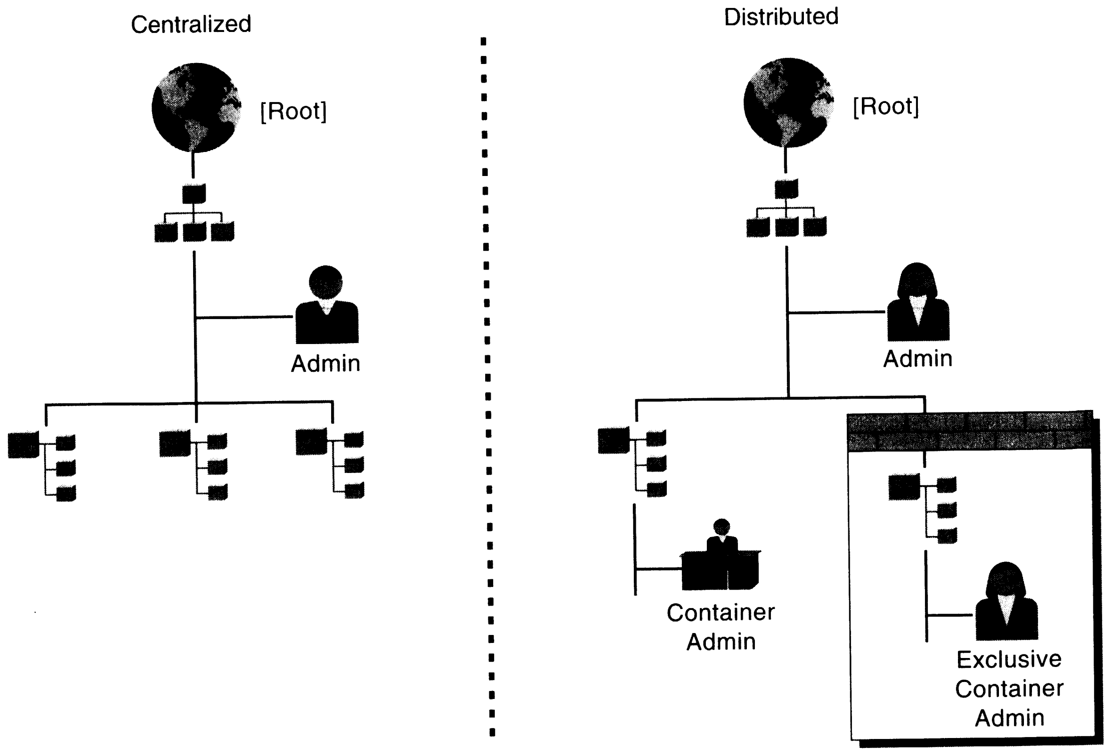

Figure 3-4

Centralized Versus Distributed Administration

Because NetWare 4 is flexible in its administration, you can design network security to meet your organization's specific needs.

NetWare 4 allows for both centralized and distributed administration. A tree can have one administrator for the entire tree, administrators for branches of the tree, or both.

Centralized Administration

Centralized administration means you have only one user or group of users with rights to the entire Directory tree. This is the NetWare 4 default.

The user Admin is created during installation and has all rights to manage the Directory tree through an explicit assignment of the [S   ] object right to [Root]. User Admin consequently inherits rights to the rest of the tree unless an IRF is applied later.

Centralized administration is appropriate for organizations with small Directory trees or organizations that want to retain a central administrator.

Some tasks should be centrally administered. These include the following:

- Naming the Directory tree
- Installing the first server
- Creating the first levels of the Directory tree
- Managing partitioning and replication (discussed in Section 4, "Partitioning and Replicating NDS and Synchronizing Time")
- Managing time synchronization (discussed in Section 4, "Partitioning and Replicating NDS and Synchronizing Time")
- Assigning container administrators

Distributed Administration

Distributed administration means that designated users are given enough NDS rights to manage branches of the Directory tree. The user who is assigned the administrative tasks at the container level (a branch of the Directory tree) is often referred to as a container administrator.

Distributed administration tends to allow administrators to respond to users' needs more quickly, especially in a large network, because control of the network does not reside with one person.

The following tasks can be distributed:

- Creating user accounts
- Creating and configuring print services
- Backing up and restoring data
- Assigning file system trustees
- Installing additional servers
- Creating workgroup managers

Setting Up the Container Administrator

There are certain methods that you should use to distribute administration tasks to a container administrator or a group of container administrators. Use the following guidelines to set up container administrators:

- If only one person will administer a container, then you will only need to make the proper rights assignments to his or her User object.
- If more than one person will administer the container, then use an Organizational Role object as the container administrator.

  If you use the Organizational Role, you should consider making an additional explicit Supervisor trustee assignment to a User object as a safety precaution.

Use extreme caution in using the Security Equal To property of a User object to create more than one container administrator (this is not recommended). If the User object is deleted and all other container administrator User objects are security equivalent to the deleted User object, all security equivalent objects will lose rights derived from the security equivalence. In the case of the exclusive container administrator, this could mean being cut off from administering a branch or container of the Directory tree.

The Organizational Role object is an ideal object to use in setting up container administrators. The Organizational Role is designed to represent a particular position within an organization where the users who occupy the position may change but the position's responsibilities do not change.

When a User object is made an occupant of an Organizational Role object, the Organizational Role object appears in the User object's Security Equal To property.

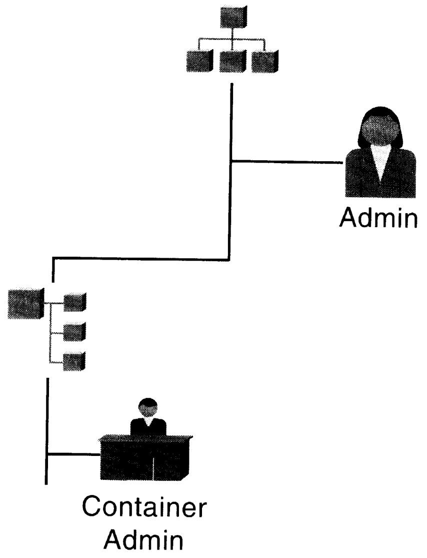

Figure 3-5: Setting Up the Container Administrator

To create container administrators using the Organizational Role object, do the following:

1. Create the Organizational Role object in the appropriate container.
2. Make the Organizational Role object a trustee of the container.
3. Assign the appropriate NDS rights to administer the container: [S] or [BCDR] or all rights [SBCDR].

   The choice of NDS rights will vary based on the network administrator's plans to implement IRFs and the container administrator's need to manage the file system in addition to the Directory tree.

4. Assign any necessary rights to the file system that pertain to a container administrator.
5. Make the appropriate User objects occupants of the Organizational Role object.

Exclusive Container Administrator

Depending on the security requirements of your organization, you may want only one user or a small group of users to have administrative rights to a particular branch.

An *exclusive container administrator* is a container administrator who has Supervisor rights in the specified container only; all rights are blocked from other network administrators higher up in the Directory tree.

For example, this assignment might be used in a government organization or in a department of a company with highly sensitive information. In this situation, you may need to cut off the rights inherited by the original Admin user.

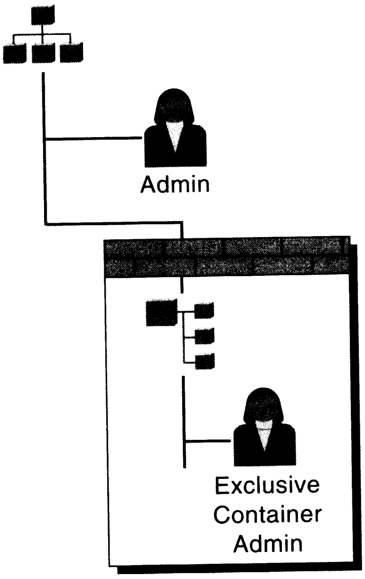

Figure 3-6: Exclusive Container Administrator

If you use the Organizational Role object to create exclusive container administrators, make an additional explicit trustee assignment to at least one of the container administrators' User objects. This will prevent a loss of administrative control of the container in the event the Organizational Role object is deleted.

To create an exclusive container administrator, you would perform the following tasks:

1. Make the container administrator's User object a trustee with the following rights assignments:

   [SBCDR] object rights and [SRCWA] property rights

Assign all rights, not just the Supervisor right. This will ensure that if the Supervisor is filtered out by an IRF, total administration is still possible.

2. Revoke inherited rights with an IRF at the workgroup container so that the original Admin does not inherit rights to the container.

   Set the IRF to [B] object rights and [R] property rights.

   Do not revoke the Browse object right; you may still want to allow other users to see that part of the tree.

3. Remove the original Admin user's trustee assignments to the container.

4. Ensure that the new container administrator has [S] object rights to himself or herself. Then remove any of the original Admin's trustee assignments to the new container administrator's User object; this will prevent the original Admin from restricting the new administrator's rights.

You must assign explicit trustee assignments that include [S] before revoking the [S] right using the IRF. If you attempt to revoke the [S] right without any explicit trustee assignments, the utilities will prevent you from doing so. However, you can still lose total control over a branch of the tree should the only user with the [S] explicit trustee assignment be deleted.

Setting Up Administrative Accounts and Assigning Rights

Administration in NetWare 4 is extremely flexible; no fixed roles for administering the tree have been defined. However, the table below offers some suggestions for administrative roles that can be useful on your network. It also discusses the types of accounts you should create for these roles.

Role	Recommended Account	Functions
Enterprise NDS Administrator	Admin User object (default Admin User object)	Install the first server. Name the Directory tree. Create the first levels of the Directory tree. Manage partitioning and replication. Manage time synchronization. Assign container administrators. Issue the initial auditor password. Upgrade servers, clients, and applications.
Container Administrator	Container Administrator/ Organizational Role object (with individual user accounts as members)	Perform data backup and restoration. Create and configure print services. Write and maintain login scripts. Monitor file server performance. Track errors. Monitor disk space usage. Assign file system trustees. Upgrade server, client, and application software.

Table 3-1: NetWare 4 Administrative Roles

Role	Recommended Account	Functions
Print Server Operator	PS Operator/ Organizational Role object (with individual user accounts as occupants)	Load and bring down the print server.
Print Queue Operator	PQ Operator/ Organizational Role object (with individual user accounts as occupants)	Delete print jobs. Change the order of print jobs. Change queue status.

Table 3-1: NetWare 4 Administrative Roles *(continued)*

The Admin account should be used only for the enterprise NDS administration functions. The individual who assumes this role should have an individual user account to use regular network services.

Even if your organization does not want a central administrator, we recommend that you keep an account that has all rights to [Root] for the sake of partitioning the Directory (discussed in Section 4, "Partitioning and Replicating NDS and Synchronizing Time"), performing NDS backups, and creating additional Organization objects.

The container administrator can create additional administrators to help with creating and deleting user accounts.

If IRFs are used to block the rights of the Admin user, ensure that each container has an administrator who can restore Admin's rights for the sake of partitioning the Directory tree (discussed in Section 4, "Partitioning and Replicating NDS and Synchronizing Time"). Creating a Container Admin Organizational Role object makes this reassignment very easy because it only involves adding Admin as a member of the Organizational Role object.

Rights for Administrative Roles

You must assign the proper rights to the administrative roles you create in your network. Proper assignment of administrative rights will help you maintain control of the network resources and minimize security breaches.

Assigning Rights

Following are the rights assignments required for the administrative roles described in the previous subsection.

Account	Assignment
Admin (User)	[S] object rights to [Root] at installation (by default)
Container Admin (Organizational Role)	[SBCDR] object rights to respective container
Print Server Operator (Organizational Role)	Added to Print Server Operator property of the respective print server
Print Queue Operator (Organizational Role)	Added to Print Queue Operator property of the respective print queue

Table 3-2: Rights for Administrative Roles

Through the assignment above, the container administrator receives enough rights for NDS and file system management. The Supervisor object right is inherited for all objects, including Server objects. Users with this right receive file system rights as well as NDS rights. This is the *only* instance when NDS security affects file system security.

Considerations for Assigning Rights

Keep the following suggestions in mind when you assign rights.

- When you assign rights to a container administrator, consider assigning all rights, not just the [S] object right. This ensures that the container administrator can still manage that branch if the [S] is blocked by an IRF at a subsequent level.
- Because an IRF can ultimately restrict the Admin user, consider adding Admin as a member of each container administrator Organizational Role. This helps prevent the loss of the Admin account's rights and ensures that the Admin user can perform all NDS functions, even if an IRF is in place.
- Decide whether this container administrator will manage the file system:
  - If yes, then grant the container administrator the Supervisor right to the Server object (if you assign the Supervisor object right at the container level, then the container administrator already has the Supervisor object right to all objects in the container).
  - If no, then create an IRF that blocks the Supervisor right to the Server object. Assign this responsibility and the appropriate file system rights to another user.

If the branch of the tree has not been created, consider giving the container administrator only the Create object right. The container administrator will then receive the [S] Supervisor object right to every object he or she creates.

Hands-On Exercise 3-2: Making Rights Assignments for a Container Administrator

Activity: Create container administrators with an Organizational Role and add three User objects from your container as occupants to create container administrators. Create a User object and add it as an occupant of the Organizational Role, also give the new User object an explicit Supervisor object right assignment to the Organizational Unit. Remove the container administrators' rights as Supervisor to the file system using a secondary trustee assignment and an IRF.

Procedure: Using NetWare Administrator, make additional rights assignments to set up a User object as a container administrator.

1. Identify two User objects from your container that you will make container administrators to your Organizational Unit.
2. Create an Organizational Role object in your container and name it Admin_*nnn*.
   a. Click on your Organizational Unit and select **Create** from the Object menu or the alternate mouse button menu.
   b. Select the **Organizational Role** icon.
   c. Click on **OK**.
   d. In the Organizational Role Name field, type

      Admin_*nnn*

   e. Click on **Create**.
3. Create a new User object and assign your name.

4. Make the two User objects you chose in Step 1 and your User object occupants of the Admin_*nnn* Organizational Role object.

   a. Click on your **Admin_*nnn*** Organizational Role object.

   b. Select **Details** from the Object menu or the alternate mouse button menu.

   c. Click on the **Select Object** button to the right of the Occupant field.

   d. Click on the **Add** button.

   e. Select one of the User objects you want to make a container administrator.

   f. Click on **OK**.

   g. Repeat Steps 3d through 3f to add the other User object and your User object as occupants.

   h. Click on **OK**.

5. Drag and drop the Admin_*nnn* Organizational Role object onto your Organizational Unit object.

6. With the Admin_*nnn* Organizational Role object highlighted in the Trustees list, assign the Supervisor object right to the Admin_*nnn* Organizational Role object.

   a. What property rights does the Supervisor object right give the Admin_*nnn* Organizational Role object?

   b. Since the Server object E*nnn* is in your Organizational Unit container object, what object rights does the Admin_*nnn* Organizational Role object have to the Server object E*nnn*? Why does the object receive these rights?

   c. Does the answer to the previous question have any implications for file system security? If so, what implications?

7. To prevent your administrative rights from being filtered, add your User object to the trustee list of the Organizational Unit object.

8. Give your User object all object rights [SBCDR]. This ensures that the container has at least one administrator with an explicit rights assignment.

9. Check the effective rights of a container administrator's User object in the PUBLIC directory. This will identify if the User object has the Supervisor object right to the Server object. What are the User object's rights in this directory?

   a. Double-click on the Volume object in your container to see the directories in the file system.

   b. Click on the **PUBLIC** directory folder.

   c. Select **Details** from the Object menu or the alternate mouse button menu.

   d. Click on the **Trustees of this Directory** page button.

   e. Click on the **Effective Rights** button.

   f. Click on the **Browse** button to the right of the Trustee field.

   g. Locate and click on the User object in the Objects field.

   h. Click on **OK**.

   i. After finding the effective rights, exit the windows relating to the PUBLIC directory.

As in the file system, you can stop inheritance of rights in NDS by making a new trustee assignment at a lower level or by using an IRF. In the following steps, you will use both methods to prevent rights inheritance. In this case, you will prevent the flow of the Supervisor (NDS) object right to the Server object.

10. Make Admin_*nnn* a trustee of the Server object.

11. Grant the Admin_*nnn* Organizational Role object Compare and Read [C R] property rights to the Object Trustees (ACL) property of the Server object.

12. What did this explicit rights assignment do to the container administrators' (members of Admin_*nnn*) User object rights in the file system?

13. Check a User object's effective rights to the Server object in your container. This User object should be a member of Admin_*nnn*.

    a. Click on the Server object in your container.

    b. Select the **Trustees of this Object** window from the Object menu or the alternate mouse button menu.

    c. Click on the **Effective Rights** button.

    d. Click on the **Select Object** button to the right of the Object Name field.

    e. Locate and click on a container administrator's User object in the Objects field.

    f. Click on **OK**.

14. To see the effect of the IRF in the following steps, remove the trustee assignment you created in Steps 10 and 11.

    a. Click on the Server object in your container.

    b. Access the **Trustees of this Object** window from the Object menu or the alternate mouse button menu.

    c. Click on the **Admin_*nnn*** Organizational Role object in the Trustees list.

    d. Click on the **Delete Trustee** button.

    e. Click on the **Yes** button in the dialog window that states, "Delete trustee assignment?"

    f. Click on **OK**.

15. Place an IRF to block the flow of the Supervisor object right to the Server object.

    a. Click on the Server object in your container.

    b. Access the **Trustees of this Object** window from the Object menu

    c. Click on the **Inherited Rights Filter** button.

    d. Click on the box next to the Supervisor object right to remove the "X".

You should see a message box stating, *"You cannot filter the Supervisor right because no User has explicit Supervisor object right to this object."*

16. Make the Admin User object that you are logged in as a trustee of the Server object and grant yourself the Supervisor object right.

17. To prevent the Supervisor object right from flowing down to the Server object, place an IRF at the Server object.

    a. Click on the Server object in your container.

    b. Access the **Trustees of this Object** window from the Object menu

    c. Click on the **Inherited Rights Filter** button.

    d. Click on the box next to the Supervisor object right to remove the "X". A shorter arrow with a line under it will appear next to the check box, indicating that the right is being blocked.

18. Check a User object's effective rights to the Server object in your container. This User object should be a member of Admin_*nnn*.

19. What rights does the User object have to the Server object? Will this affect the User object's rights to the file system?

Assigning Additional Rights Using the Group Object

If you have several users who need the same rights, you can add them as members of a Group object. Then you can assign rights to the Group object, reducing the number of individual rights assignments.

When a User object is made a member of a Group object, the Group object appears in the User object's Security Equal To property.

Group Object: Users in the Same Container

Since a container object acts like a natural group, you do not need to create a Group object to assign NDS rights for users in the same container. Under most circumstances, Group objects are used to grant group members rights to the file system directories based on the group members' needs.

Example: All users that need access to specialized accounting software could be assigned as members of the Accounting group and be given file system rights to a directory named Accounting.

Global Group

The Group object can also be used to give resource access to users outside of the resource's container. This type of group is referred to as a global group.

If your network involves WAN links, you will need to consider the implications of using the global group. These considerations will need to be examined closely when designing a Directory tree that spans various geographical locations.

Example: Division managers need access to a printer at the corporate headquarters to print daily reports. They also need access to a centralized location to get up-to-date product information stored on a server at corporate headquarters.

The division managers are made members of the Managers Group object, which resides in the corporate container. The Managers Group object is then given rights to access the proper directory where the product information is stored and to print to the printer.

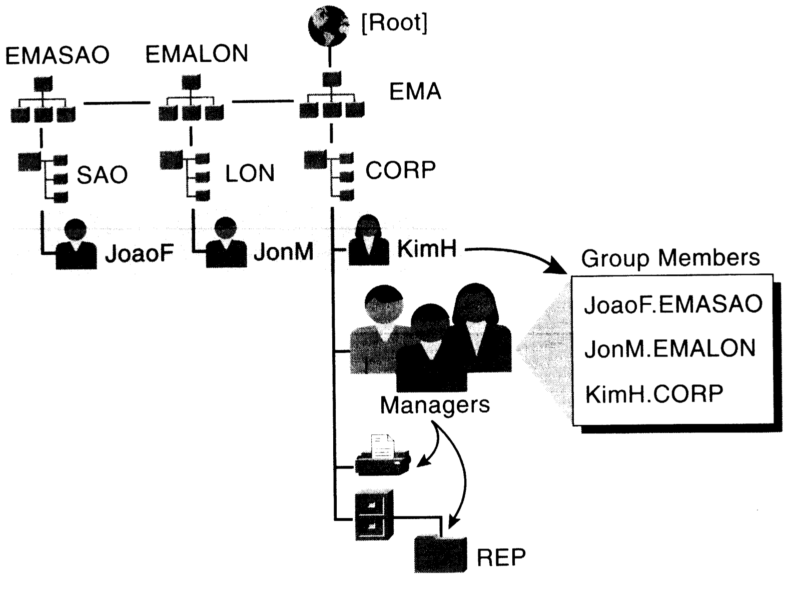

Figure 3-7: Global Group

Objects Requiring Additional Rights

Certain objects in the Directory tree require additional NDS rights assignments in order to access a particular network resource. The required rights assignments occur when the object is found in a different container than the User object accessing the resource. The following objects may require additional rights assignments:

- Profile Login Script
- Directory Map

Profile Login Scripts

To run a profile login script, a user needs the Read property right to the Login Script property of the Profile object. If the User and Profile objects are in the same container, the User object does not require additional rights; he or she already has the Read property right to any object in the container.

The following steps explain how to assign the necessary rights to run a profile login script:

1. Make each user who will run the script a trustee of the Profile Login Script object; the Browse object right will be granted by default.
2. Grant each user the Read property right to the Login Script property.

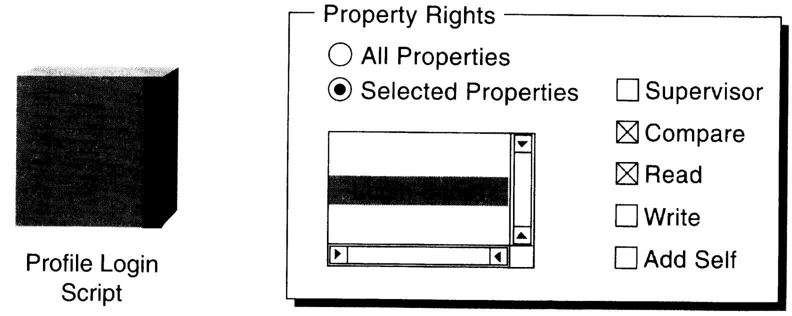

Property Rights

○ All Properties

● Selected Properties

□ Supervisor

☒ Compare

☒ Read

□ Write

□ Add Self

Figure 3-8: Assignment of Profile Login Script

Directory Map Objects

The Directory Map object reduces redundant network administration. A Directory Map points to a directory. When the location of the directory is moved, you can update the Directory Map object to reflect the change instead of going to all the user workstations and changing all the network drive mappings.

A user must have the necessary rights to read the properties that point to the proper directory. If the User object and the Directory Map object reside in the same container, no additional rights assignments are required. If the Directory Map object is placed outside the User object's container, you will need to make the user a trustee of the Directory Map object and grant the Read and Compare property rights with the All Properties option. (You can also make the appropriate rights assignment through the Selected Properties option by granting the Read and Compare property rights to the Path property.)

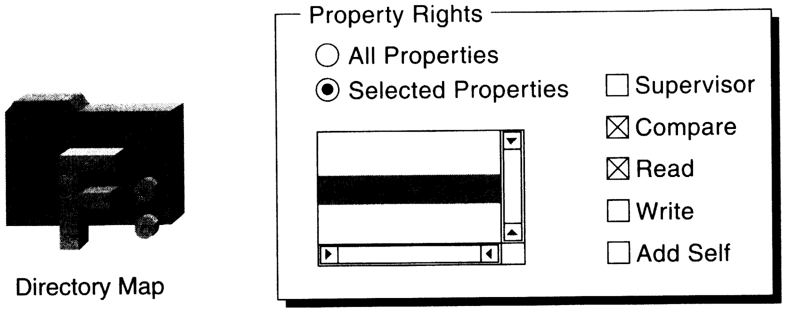

Property Rights

○ All Properties

● Selected Properties

□ Supervisor

☒ Compare

☒ Read

□ Write

□ Add Self

Figure 3-9: Directory Map Objects

Additional NDS Rights for Special Users

The following is a description of two types of users, mailing list administrators and traveling users, that require specific rights assignments. Both of these special users could be established with either a Group object or an Organizational Role object placed in the proper container.

Mailing List Administrator

Some users, such as mailing list administrators, have to manage certain network resources.

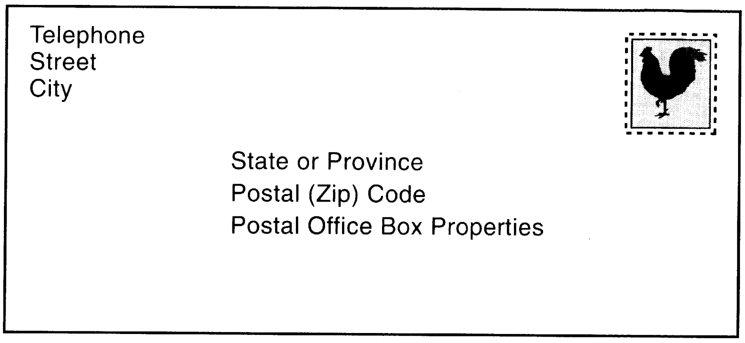

Figure 3-10: Mailing List Administrator

In setting up a mailing list administrator, you will need to make him or her a trustee of each User object and grant Read and Write property rights to the following Selected Properties of each User object:

- Telephone
- Street
- City
- State or Province
- Postal (Zip) Code
- Postal Office Box properties

Rights should not be assigned to the mailing list administrator using the All Properties rights selection. This would provide too many rights, including the right to the Object Trustees (ACL) property.

Traveling Users

Planning for the traveling user or a temporary user will require additional NDS rights and file system rights assignments. Various types of traveling users exist. For this discussion, we will define the traveling user as follows:

- A user who divides his or her time between two offices
- A user who is in a location temporarily (Example: Someone working on a specific project that requires a few days to a few weeks)
- A user who travels to various locations on a regular basis (Example: Regional sales manager who travels to regional sales offices)

Figure 3-11: Traveling Users

You should consider the following issues when working with traveling users:

- Access to applications
- File storage in a directory
- Access to files in another location
- Access to resources such as printers and E-mail
- Authenticating to NDS and to NDS objects
- Type of computer (notebook or desktop)
- Number of traveling users (especially the users whose objects reside in the same location or container)

After considering these issues, you can plan and implement a method for the different types of traveling users. Because of the flexibility of NetWare 4, each type of traveling user can receive his or her needed resources in different ways.

The following examples should help you plan network security and access for the traveling user. Each one of these traveling user scenarios could be handled with more than one method.

Traveling User — Time Between Two Offices

The user who divides his or her time between two offices needs similar resources in two different locations. The user needs to access the same applications, print to a printer, and store files. This traveling user situation could be handled with the following:

- Create two User objects, one in each location.
- Give appropriate rights to any resources needed in each location.
- Give appropriate rights to read any Profile Login Script objects and Directory Map objects.

Traveling User — On Location Temporarily

The user who is at a location temporarily needs to access the appropriate directories and use printers. A possible solution to consider might be the following:

- Create an Alias object or assign the user to a Group object or Organizational Role object.
- Assign the appropriate rights to the Alias, Group, or Organizational Role object for directories and network resources.
- Give appropriate rights to the Alias, Group, or Organizational Role object to read any Profile Script objects and Directory Map objects.

Traveling User — Travels to Various Locations

The user who travels to various locations on a regular basis requires many of the same resources as well as access to applications and files in the file system on the network. You can use several methods to handle this type of traveling user, including the following:

- Create an Alias object or assign the user to a Group object or Organizational Role object.
- Assign the appropriate rights to the Alias, Group, or Organizational Role object for directories and network resources.
- Give appropriate rights to the Alias, Group, or Organizational Role object to read any Profile Script objects and Directory Map objects.

Troubleshooting NDS Rights

Most troubleshooting with NDS rights involves finding users who have unauthorized access to objects and their properties and determining why users cannot access the proper network resources.

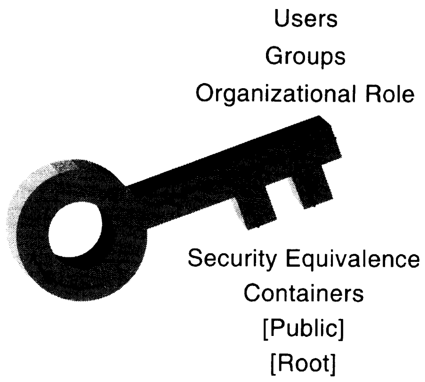

Figure 3-12: Troubleshooting NDS Rights

Unauthorized Access

If someone has more access to a resource than he or she should have, you must determine where the rights for the resource are coming from.

The best method to discover unwanted rights for a user is determining the user's effective rights. To determine a user's effective rights, use NetWare Administrator to do the following:

1. Click on the resource you want to check.
2. Access **Trustees of this Object** from the Object menu.
3. Click on the **Effective Rights** button.
4. Click on the **Browse** button to the right of the Object Name field.
5. Click on the User object in the Objects list.
6. Click on **OK**.

You may have to walk up and down the tree to find the effective rights for a particular user and identify where the rights are coming from. Once you identify where the rights begin, you can determine how the user received the rights, as explained below:

1. Check explicit trustee assignments for the following:
   - User object
   - Groups and Organizational Role the user is a member of
   - Security Equivalences the user may have
   - Containers the user is in up the tree to the [Root]
   - Rights given to the [Public] trustee
   - Rights given to the [Root] of the tree
2. Check the inherited rights for the items listed in Step 1.
3. Access the **Trustees of this Object** window to check the trustee assignments made to that container using NetWare Administrator.
4. For a directory, use NetWare Administrator to perform the following:
   a. Double-click on the **Volume** object.
   b. Click on the directory for which you want to check a user's rights.
   c. Select **Details** from the Object menu.
   d. Click on the **Trustees of this Directory** page button.
   e. Click on the **Effective Rights** button.
   f. Click on the **Browse** button to the right of the Trustee field.
   g. Click on the desired User object.
   h. Click on **OK**.
5. Determine if the user is in, or has security equivalence to, containers or groups that have rights to the container or directory.

Unable to Access Resource

To determine why a user may not be receiving rights to a container or directory, ask the following questions.

Group/Organizational Role:

Has the user been made a member/occupant?

Security Equivalence:

Has the user been made security equivalent?

Container:

Has the container been assigned rights?

Are you sure that everyone in the container should receive the same rights? If yes, make sure that the assignment is made at the container level. If no, create a group and assign rights to the group. Then place the users who need the special rights in the group.

User:

Has the user been assigned the rights?

Is an IRF placed at the container or directory in question?

- If yes, make an explicit trustee assignment to the container or directory.
- If no, check the parent containers for IRFs and then make explicit assignments.

Hands-On Exercise 3-3: Troubleshooting NDS Security

Activity: Using the following scenario, troubleshoot NDS rights.

Scenario: You are in the \GROUPS\ACCNT\DB directory assigning file system rights to MarcJ and RussC. You decide to check their effective rights and you find that both users have too many rights.

Procedure: Use NetWare Administrator to troubleshoot the above scenario. Find out the source of MarcJ and RussC's rights.

1. What are MarcJ's effective file system rights to \GROUPS\ACCNT\DB?
2. Walk up the file system directory, checking effective rights for MarcJ. Is MarcJ receiving his rights from somewhere in the file system?
3. Check NDS rights, starting in the container with the Server object.
4. What object in the container would give MarcJ rights to the file system? Is he a trustee for that object?
5. Check NDS rights at the container and parent containers.
6. Where does MarcJ get his rights?
7. Perform Steps 1 through 6 for RussC.
8. If none of the above steps identify where RussC receives his rights, in what other ways can RussC get his rights?
9. Where does RussC get his rights from?
10. What is the appropriate action to remove MarcJ's rights to the file system and still allow him to have the same object rights in the container EMACORP?
11. What is the appropriate action to remove RussC's ability to gain rights to the file system, except through an explicit assignment?

Enabling Network Auditing

Another method of ensuring network security is to allow independent auditors to audit network events. NetWare 4 allows you to enable auditing to track events in NDS and in the file system. After the auditing has been performed, you will need to disable the network auditing.

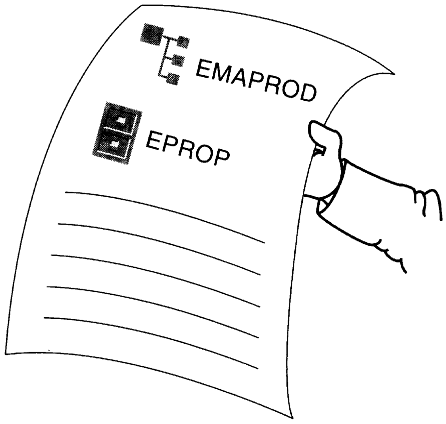

Figure 3-13: Enabling Network Auditing

As a network administrator, you will need to do the following:

- Create a User object for the auditor if one has not already been created.
- Enable network auditing for an auditor.
- Create a current password.
- Give the password to the auditor (he or she will then change the password to establish independence from the network administrator).
- When the auditing is complete, disable auditing on both containers and volumes.

When auditing is enabled for an NDS container, it is enabled for that container only; it is not enabled for subordinate containers. Likewise, when auditing is enabled for a volume, it is enabled for that volume only; this reduces overhead on other volumes.

At the NDS container level, events relate to the use of NDS. For example, you can audit the NDS container level to monitor the creation of User objects.

At the volume level, events relate to the use of files and directories, queues, or NetWare servers. For example, you could audit the volume level when you want to track the number of times a certain user opens a specific file.

Enabling Auditing for an NDS Container

To enable auditing for an NDS container, do the following:

1. At the DOS prompt, type the following:

   AUDITCON <Enter>

   The Available Audit Options menu appears. The current NetWare server and volume appear at the top of the screen.

2. Select **Audit Directory Services**.

   The Audit Directory Services menu appears. The session context appears at the top of the screen.

3. (Optional) Change the context:

   a. Select **Change Session Context**.

   b. At the prompt, type the context of the container for which auditing is being enabled; then press <Enter>.

      The Audit Directory Services menu appears.

4. Select **Audit Directory Tree**.

   A list of containers appears.

5. Highlight the container to be audited; then press <F10>.

   The Available Audit Options menu appears.

6. Select **Enable Container Auditing**.

7. At the prompt, type an auditor's password for the container; then press <Enter>.

8. At the prompt, retype the password; then press <Enter>.

9. Notify the auditor of the password.

For further information, see "Auditing NetWork Events" in *Supervising the Network.*

Enabling Auditing for a Volume

The procedure for enabling auditing at the volume level is similar to enabling auditing at the NDS container level. However, menu selections pertain to the volume being audited.

To enable auditing for a volume, do the following:

1. At the DOS prompt, type the following:

   AUDITCON <Enter>

   The Available Audit Options menu appears. The current NetWare server and volume appear at the top of the screen. Change to the volume you want to audit.

2. Select **Enable Volume Auditing**.

3. If more than one volume is located on the Server in the current context, highlight the volume to be audited; then press <F10>.

4. At the prompt, type an auditor's password for the volume; then press <Enter>.

5. At the prompt, retype the password; then press <Enter>.

6. Notify the auditor of the password.

Disabling Auditing for an NDS Container and Volume

Once the auditor has completed the audit and has prepared his or her reports, you will need to disable the auditing for the container and volume for which the auditing occurred.

Disabling Auditing for an NDS Container

1. At the DOS prompt, type the following:

   AUDITCON <Enter>

   The Available Audit Options menu appears. The current NetWare server and volume appear at the top of the screen.

2. Select **Audit Directory Services**.

   The Audit Directory Services menu appears. The session context appears at the top of the screen.

3. (Optional) Change the context:

   a. Select **Change Session Context**.

   b. At the prompt, type the context of the container for which auditing is being disabled; then press <Enter>.

      The Audit Directory Services menu appears.

4. Select **Audit Directory Tree**.

   A list of containers appears.

5. Highlight the container for which auditing is being disabled; then press <F10>.

6. With the **Auditor Container Login** highlighted, press <Enter>.

7. Enter the auditor password and press <Enter>.

   The Available Audit Options menu appears.

8. Select **Auditing Configuration** and press <Enter>.

9. Select **Disable Container Auditing** and press <Enter>.

10. Select **Yes** and press <Enter>.

11. Press <**Escape**> until the “Exit?” prompt appears.

12. Select **Yes** and press <Enter>.

Disabling Auditing on a Volume

1. At the DOS prompt, type the following:

   AUDITCON <Enter>

   The Available Audit Options menu appears. The current NetWare server and volume appear at the top of the screen. Change to the volume for which you are disabling auditing.

2. Select **Auditor Volume Login** and press <Enter>.

3. Enter the auditor's password (you may need the password from the auditor to perform this step).

   The Available Audit Options menu will appear.

4. Select **Auditing Configuration** and press <Enter>.

5. Select **Disable Volume Auditing** and press <Enter>.

6. Select **Yes** and press <Enter>.

7. Press <**Escape**> until the "Exit?" prompt appears.

8. Select **Yes** and press <Enter>.

Hands-On Exercise 3-4: Enabling and Disabling Network Auditing

Activity: Enable and disable auditing for your container EMA*nnn* and for the volume E*nnn.*

Procedure: Use the steps described previously to complete the exercise.

1. Enable auditing for the NDS container EMA*nnn.*
2. For the password, type EMA*nnn.*
3. Enable auditing for the Volume E*nnn*_SYS.
4. For the password, type E*nnn.*
5. Disable auditing for the NDS container EMA*nnn.*
6. Disable auditing for the Volume E*nnn.*
7. How does the auditor ensure that he is isolated from the network administrator?

Summary

In addition to understanding the components of NDS security, you need to know how to implement NDS security to distribute network administration, set up special users, and initiate network auditing.

In the case of a user not having enough rights or too many rights, you need to be able to troubleshoot the problem and implement a solution to correct the problem.

Notes

SECTION 4 Partitioning and Replicating NDS and Synchronizing Time

Introduction

In this section, you will partition the Directory, replicate it, and place the replicas on multiple servers across the network. These tasks will help you secure and enhance your network. You will also learn to synchronize the time among servers so that time-critical services function properly.

Objectives

After completing this section, you will be able to do the following:

1. Partition the Directory to allow proper replication and distribution to multiple servers.
2. Place replicas in a way that fosters network accessibility, performance, and fault tolerance.
3. Anticipate what partitions and replicas are created by default when a server is installed in a Directory tree.
4. Move a container object and its subordinate objects within the Directory tree.
5. Configure servers to synchronize time in a NetWare 4 network.
6. Given a scenario that includes a WAN topology, choose between Service Advertising Protocol (SAP) and specified time sources to synchronize time among servers.

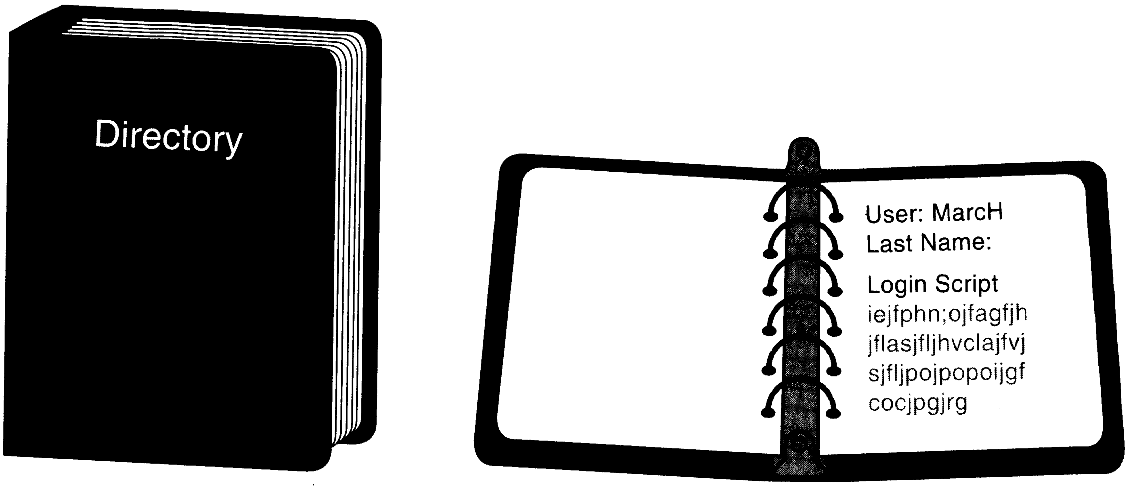

Information on all objects
in the Directory tree

Figure 4-1

The Directory

NDS

The Directory has the following characteristics:

- It is a database that replaces the bindery.
- It contains data on all objects in the Directory tree, including objects' names, rights, and property values.
- NDS uses the Directory for access control (checking to see whether an object has rights to perform an action in NDS).
- NDS uses the Directory for authentication, an important part of login.
- Except for the Server and Volume objects, the Directory does *not* contain information on the file system.

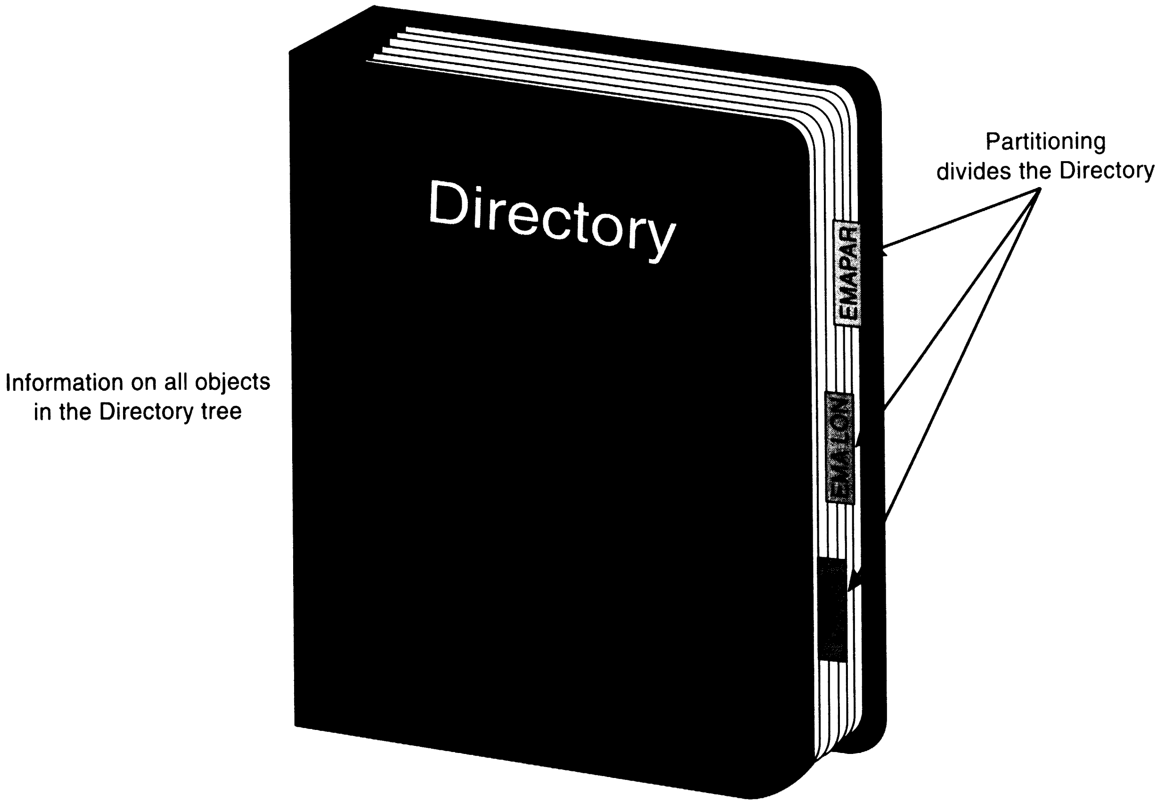

Information on all objects in the Directory tree

Partitioning divides the Directory

Figure 4-2

Partitioning

Partitioning is the process of dividing the Directory. Since a large Directory can contain information on thousands of objects, partitioning makes it possible to apportion the Directory among multiple servers. Partitions are managed by the network administrator.

Partition Root Container

Partitioning occurs along the boundaries of container objects, often including more than one container. All the leaf objects in a container are in the same partition as the container.

A partition is named by the container closest to the [Root]. This container is called the *partition root*, as illustrated in the following figure.

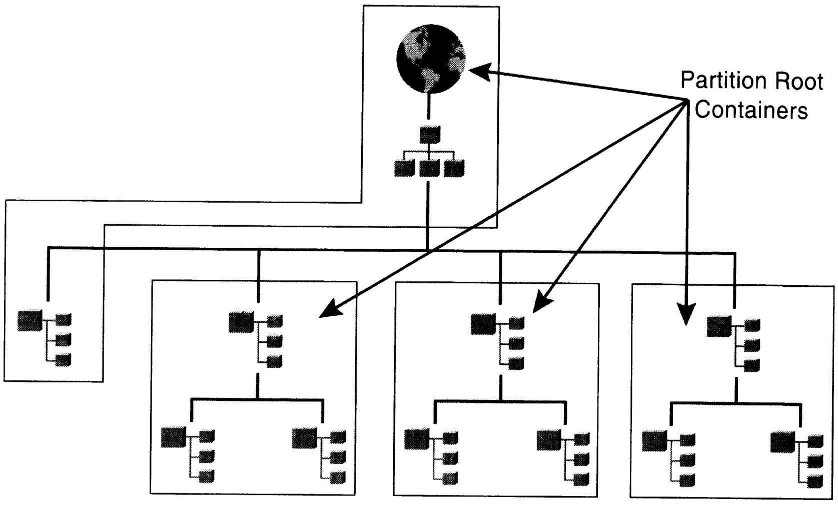

Figure 4-3: Partition Root Container

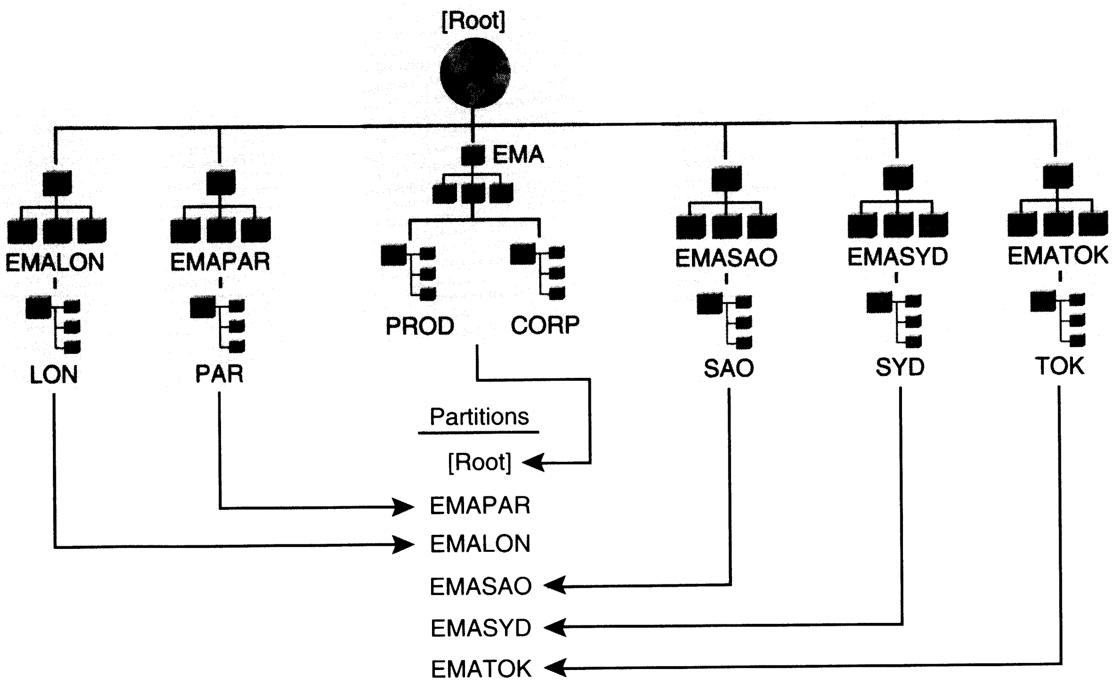

Figure 4-4

Partitioning the EMA Directory Tree

Figure 4-4 shows the partitioning of the EMA Directory tree. Note that the partitions are named after their root containers.

Parent and Child Partitions

A partition that exists immediately above the root of another partition is called that partition's parent.

A partition that exists immediately below another partition is called that partition's child, as shown in the figure below.

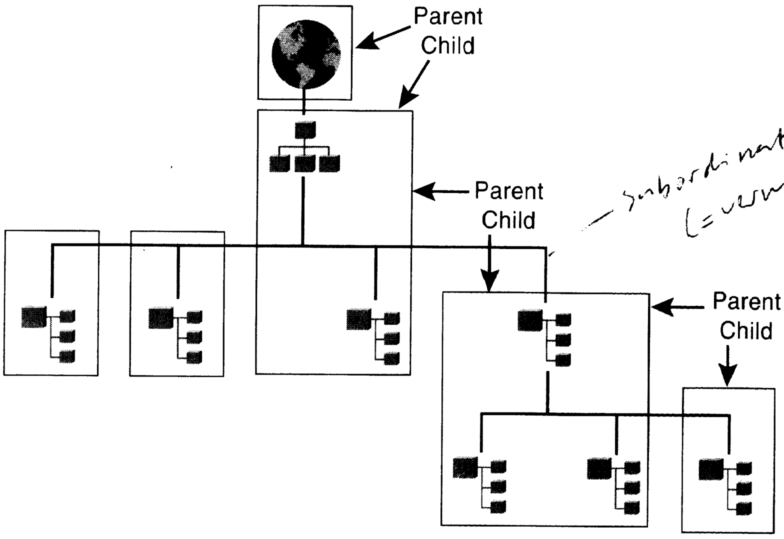

subordinate
(= [illegible])

Figure 4-5: Parent and Child Partitions

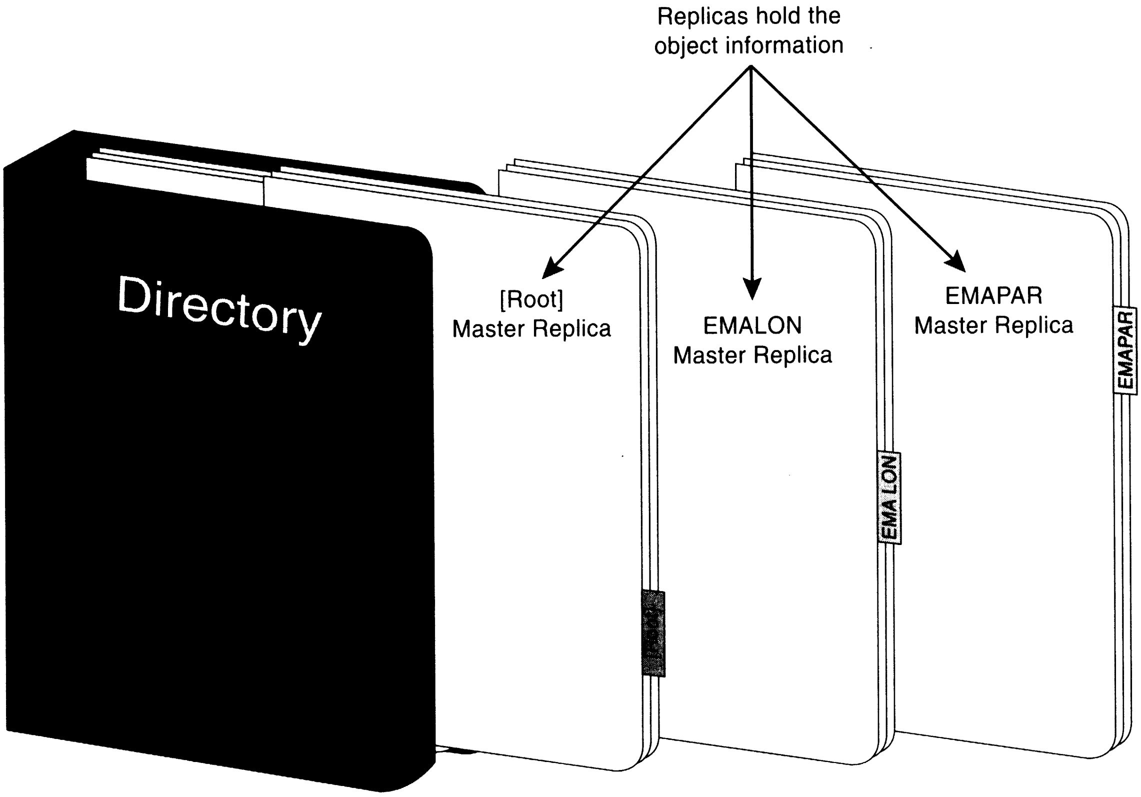

Figure 4-6

Replicas

Replicas contain the actual Directory data for objects within the boundaries of a partition. Replicas are stored on NetWare 4 servers around the network and are managed by the network administrator.

All the replicas in the Directory tree compose the Directory.

Types of Replicas

NDS uses four types of replicas. These replicas, listed below, have different functions in the Directory:

- Master replicas
- Read/write replicas
- Read-only replicas
- Subordinate references

Master Replicas

When you first define a partition, an original *master replica* is created. Each partition has only one master replica.

When you change an object contained in a master replica, that change will be propagated to all other replicas of that partition.

When you redefine a partition boundary by splitting a partition or joining it with another, you must be able to access the server that holds the master replica of the partition.

Master replicas can be used for authentication (part of login).

Read/Write Replicas

A *read/write replica* is a copy of the partition. You can have multiple read/write replicas for a partition.

When you change objects in a read/write replica, those changes will be propagated to all other replicas of the partition. However, you cannot use read/write replicas to redefine the partition boundaries.

If the server that holds the original master replica goes down, you can make a read/write or read-only replica of the master. If the original master replica ever comes back online, it will be immediately deleted; thus, only one master replica will exist at any time.

Read/write replicas can be used for authentication (part of login).

Read-Only Replicas

A *read-only replica* receives and propagates changes made to master and read/write replicas. You can have multiple read-only replicas for a partition.

Read-only replicas are used for searching and viewing objects; you cannot change objects on a read-only replica.

Read-only replicas do not support user authentication.

Subordinate References

Subordinate references are very different from the other replica types. They do not contain object data, but they point to the replica that does.

Subordinate references are maintained by NDS, not by the administrator. NDS uses the *subordinate reference* replica type to facilitate tree connectivity ("walking" the tree).

A subordinate reference is created automatically on a server when the server contains a replica of a partition but not of that partition's child.

Stated another way, subordinate references are created on servers "where the parent is, but the child is not."

If you add a replica of that child partition to the server, the subordinate reference is automatically removed.

Subordinate references do not support user authentication or viewing or managing objects in the partition. Instead, they refer to the read/write or master replica that can support these services.

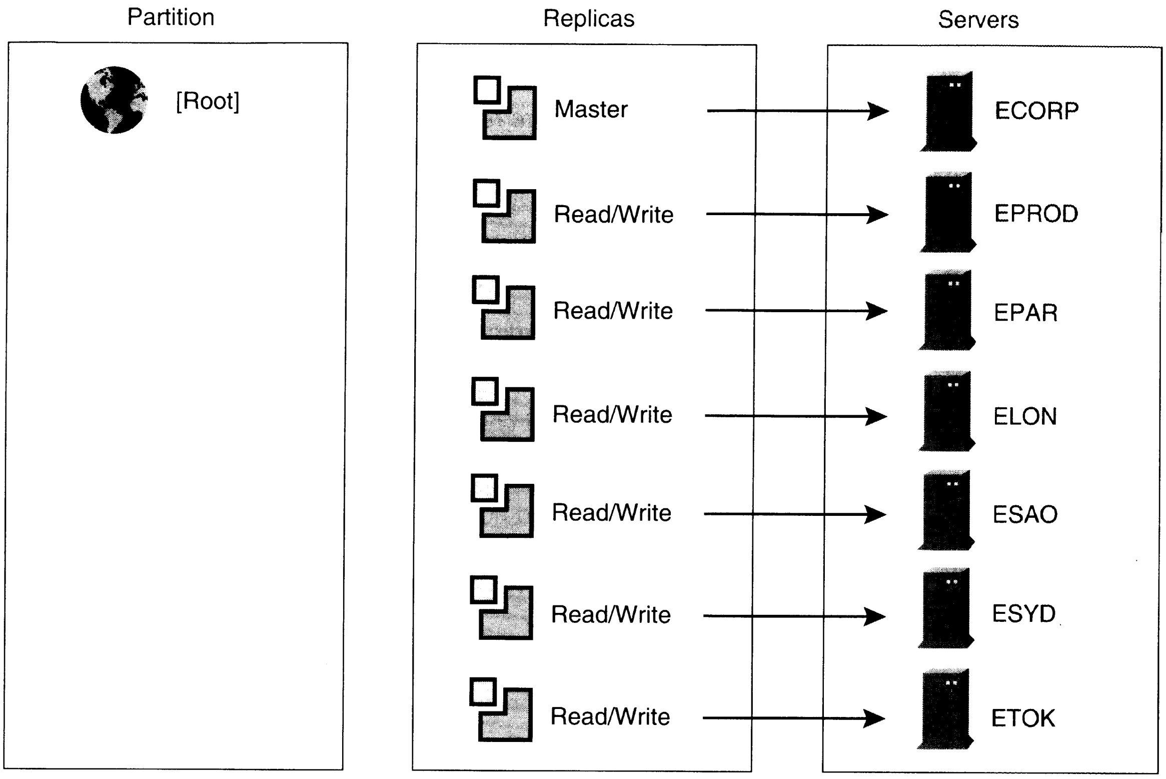

Figure 4-7

Fault Tolerance with Replication

Creating multiple NDS partitions does not, by itself, increase fault tolerance or improve performance of the Directory. However, strategically using multiple replicas does.

- Replication provides fault tolerance for the Directory by redundantly storing the Directory data on various servers.
- Replication can improve response time and reduce network traffic when services (such as LOGIN) are requested.

Replica Ring

A replica ring includes all servers with replicas of the same partition. The replica ring for the [Root] partition of the EMA Directory tree is shown in Figure 4-7.

Synchronization Cost of Multiple Replicas

When a change is made to an object, that change is communicated to all replicas in the replica ring. The more replicas in a replica ring, the more communication is required to synchronize changes.

The time cost of synchronization is greater when the servers keeping a replica ring are separated by relatively slow WAN links.

Therefore, the limiting factor in creating multiple replicas is the amount of processing time and traffic required to synchronize. Novell recommends three or more replicas in a ring, depending on how much synchronization cost the network can accommodate.

Working with Partitions and Replicas

This subsection describes the defaults for partitioning and discusses guidelines for managing partitions and replicas.

Defaults for Partitions and Replicas

When you install a new NetWare 4 server, the partition into which you install it expands. No new partition is created. Replication of the partition depends on how many servers already exist in the partition.

First Server Installed in a Directory Tree

The first partition is created at the time of installation of the first NetWare 4 server. A master replica is stored on that server's SYS: volume.

Subsequent Servers in the Directory Tree

For subsequent server installations, the following default partitioning and replicating rules apply:

- When a new server is installed in an existing Directory tree, the new server becomes part of the existing partition, even if new container objects are created. No new partitions are created.
- The second and third new servers in the partition receive a read/write replica of the partition. The fourth and subsequent servers do not receive any replicas.

Merging Directory Trees

When you merge Directory trees, servers in the source tree that contain replicas of the [Root] partition receive a read/write replica of the new [Root] partition. They also receive subordinate references to the [Root] partition's child partitions.

Servers in the target tree that contain replicas of the [Root] partition receive subordinate references of the top-level partition of the source Directory tree.

We suggest that you evaluate partitions and replicas after merging Directory trees to reduce the number of subordinate references. See Table 4-1 (on page 4-8) for a sample of post-merger partitions and replicas.

Upgrading Bindery Servers

A NetWare 3 server upgraded to NetWare 4 receives a read/write replica of all partitions containing the server's bindery contexts. Bindery context will be covered in Section 6, "Integrating and Managing NetWare 3."

Guidelines for Managing Partitions and Replicas

A partition should be replicated for two reasons: to create fault tolerance and to increase efficiency for users. Use these guidelines to help you plan fault tolerance without increasing your network traffic:

- To meet fault tolerance needs, plan for three or more strategically placed replicas of each partition. Additional replicas can be created to make login and access time faster for users.
- Create partitions that follow the boundaries of each workgroup and its associated resources. Place replicas of each partition on servers that are physically close to the workgroup that uses the information in that partition.
- If your network requires bindery services, make sure each server contains a master or read/write replica that contains the bindery context (the context set by the SET BINDERY CONTEXT statement in the AUTEXEC.NCF file).

For the most part, following these guidelines and using some networking common sense will make you successful at managing partitions and replicas.

Replicating the [Root] Partition

Be sure to replicate the partition that includes the [Root] object. This is the most important partition of the tree; if you lose this partition, your Directory tree becomes inaccessible.

However, you should be careful of replicating the [Root] or other high-level partitions too many times. Since these partitions often have multiple child partitions, each server that receives replicas of the [Root] also receives subordinate references for all its child partitions. This may create more synchronization cost than you intended.

Managing Subordinate References

Any partition operation, such as creating, merging, or deleting a partition, affects subordinate references as it would any other partition. Therefore, you need to consider subordinate references in your partition and replica design.

For example, you would not want to have subordinate references linked across unreliable WAN connections to other replicas of the same partition. If you wanted to make a partition change, such as merging partitions and the link was not available, you would not be able to complete the action until the subordinate references across that link became available. In the meantime, the Directory would not be synchronized.

We recommend the following steps for reducing the number of subordinate references:

- **Create few partitions at the top levels of the tree.** Such a design creates fewer subordinate references than a top-heavy tree that has more partitions at the top than at the bottom.
- **Avoid creating unnecessary partitions.** Workgroup boundaries generally determine the number of partitions required in any tree. You should partition your tree according to how network resources are used and where they are located in the Directory tree.
- **Minimize the number of replicas of the [Root] partition and other parent partitions**. Remember that every server that contains a replica of a parent partition receives subordinate references of that parent's child partitions.

Exercise 4-1: Viewing Partitions and Replicas

In this exercise, you will see how NetWare Administrator displays partitioning for the EMACORP tree.

Procedure: Write your answers in the space provided.

1. At your workstation, log in as administrator for your organization. Type

   LOGIN E*nnn*/.ADMIN.EMA*nnn* <Enter>

2. Change your current context to [Root]. Type

   CX [Root] <Enter>

3. Launch MS Windows and NetWare Administrator.
4. From the Tools menu, select **Partition Manager**.
5. What do the icons next to the word "[Root]" mean? (Click on the **Help** button for more information.)
6. Click on the **Replicas...** button.
7. On what servers are the replicas of this partition stored?
8. What types of replicas are they?
9. Close the **Partition Replicas** window.

10. Double-click on the **[Root]** object icon.
11. Double-click on the **EMA*nnn*** Organization object.
12. What is the difference between the CORP Organizational Unit and the EMASAO Organization, as indicated by the icons? (Click on the **Help** button for more information.)
13. Double-click on the **CORP** Organizational Unit.
14. What classes of objects are displayed?
15. Click on the **ECORP** server icon.
16. Choose the **Server Partitions...** button.
17. ECORP contains replicas of which partitions?

The EMA Replica Table

A table similar to the following can be helpful in analyzing your partitioning and replication.

[illegible]	[illegible]					
	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
ECORP	Master	Sub. ref.	Sub. ref.	Sub. ref.	Sub. ref.	Sub. ref.
ELON	Read/write	Master	Sub. ref.	Sub. ref.	Sub. ref.	Sub. ref.
EPAR	Read/write	Sub. ref.	Master	Sub. ref.	Sub. ref.	Sub. ref.
EPROD	Read/write	Sub. ref.	Sub. ref.	Sub. ref.	Sub. ref.	Sub. ref.
ESAO	Read/write	Sub. ref.	Sub. ref.	Master	Sub. ref.	Sub. ref.
ESYD	Read/write	Sub. ref.	Sub. ref.	Sub. ref.	Master	Sub. ref.
ETOK	Read/write	Sub. ref.	Sub. ref.	Sub. ref.	Sub. ref	Master

Table 4-1: EMA Replicas

Moving Container Objects

Now that you have learned how to partition and replicate the Directory, you can move a container, including all the objects below the container.

Since moving a container object directly affects partitions, it is considered a partitioning operation.

Use the Partition Manager tool in NetWare Administrator to move a container.

To move a container, you must first make it the root of a partition. If you fail to do so, you will see the following message:

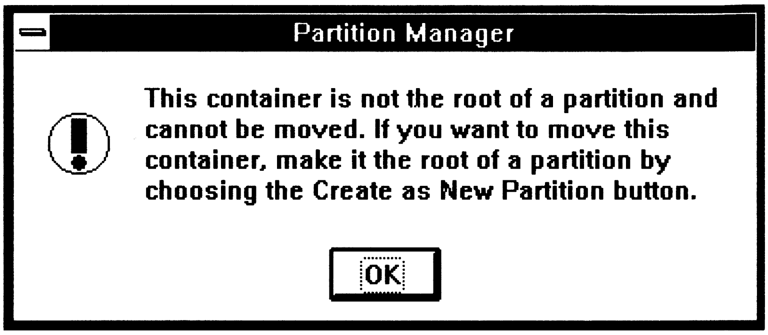

Partition Manager

This container is not the root of a partition and cannot be moved. If you want to move this container, make it the root of a partition by choosing the Create as New Partition button.

Figure 4-8: Move Container Warning

Implications of Moving Containers

Moving containers changes the NDS distinguished name for each of the objects. To adapt to the changes in the NDS structure, each user in the container that has been moved must do the following:

- Log in using his or her new NDS distinguished name.
- Change the "NAME CONTEXT =" statement in NET.CFG on the workstation.

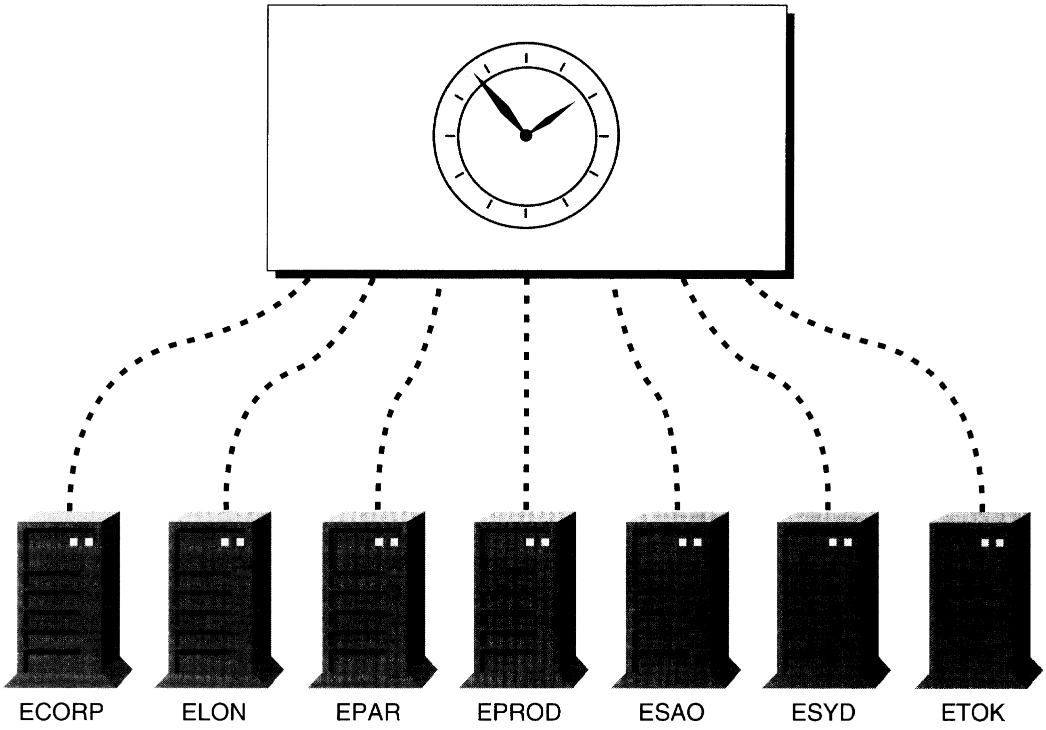

Figure 4-9

Time Synchronization

Time synchronization is a service that maintains a consistent time standard across the network servers.

- File systems use server time to apply time stamps to file and directory operations.
- Messaging applications use time stamps for messages.
- NDS uses server time to help properly collate changes to the Directory database by attaching a time stamp to Directory requests and changes.

How NDS Uses Time

In NDS, any changes you make to objects occur first in a single replica. From there, the change is communicated throughout the network to other replicas in the replica ring.

If the same object is modified in two different replicas of the database or modified twice in the same replica, the order in which the modifications were made must be preserved to ensure database integrity.

Each NDS action is therefore given a *time stamp.* Time synchronization maintains consistent time among servers so that all time stamps are accurate among servers in a replica ring.

Universal Time Coordinated (UTC)

Time stamps use Universal Time Coordinated (UTC), a time system that adjusts for the local time zone of the server. UTC corrects the local server time to get the equivalent of Greenwich Mean Time (GMT).

The formula for calculating UTC is as follows:

local time +/- time zone offset (- *daylight savings time offset*) =UTC

For instance, the Provo, Utah, time is seven hours behind GMT. Therefore, if the time in Provo is 11:00 and there is no daylight savings time, UTC time is 18:00.

When you install a server, you are prompted to set time zone offset parameters in the following screen:

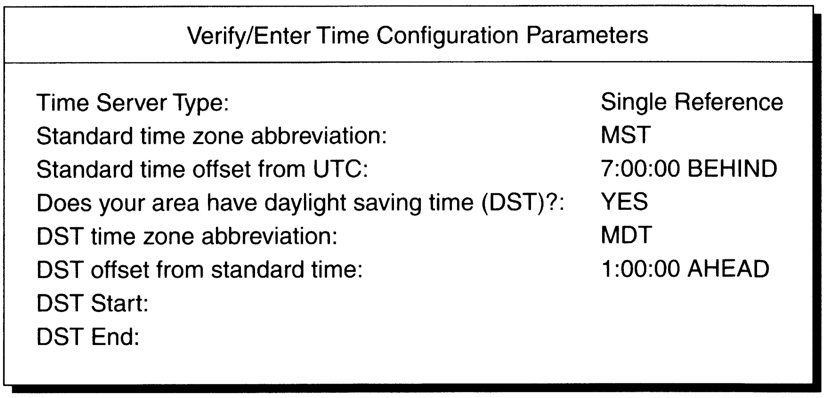

Verify/Enter Time Configuration Parameters

Time Server Type:	Single Reference
Standard time zone abbreviation:	MST
Standard time offset from UTC:	7:00:00 BEHIND
Does your area have daylight saving time (DST)?:	YES
DST time zone abbreviation:	MDT
DST offset from standard time:	1:00:00 AHEAD
DST Start:	
DST End:	

Figure 4-10: Time Synchronization Installation Parameters

Time Servers

All NetWare 4 servers are *time servers* of some kind. Time servers provide a consistent source for the time stamps that NDS uses to identify and order Directory events. Time servers ensure that the time stamps are accurate.

NetWare 4 provides four types of time servers:

- Reference servers
- Primary servers
- Single Reference servers
- Secondary servers

Secondary servers do not provide time; they adjust their clocks to one of the other time server types. Reference, Primary, and Single Reference servers are called *time providers.*

Time providers determine time by dictation or arbitration, depending on the type of time server. When servers agree on a time, they are *synchronized.*

The server types and their functions are described in the table below.

Single Reference	Yes	Default configuration. Same as Reference, but always claims to be synchronized with network time. Cannot coexist with Primary or other Reference servers.	Hardware clock or external source.	No	Secondary and clients
Primary	Yes	Polls all other time sources, votes to determine correct network time, and compensates for clock errors. Sets synchronization status based on its deviation from calculated network time, without regard to status of other time sources polled.	Reference. If no Reference server exists, you must have at least one other Primary time server.	Yes (50% correction per polling interval)	Secondary and clients
Reference	Yes	Same as Primary, but does not adjust its internal clock. Provides a central point of time control for entire network. A higher-priority time source than a Primary server, since it is considered more reliable. Reliability may be provided through commercial products such as radio clocks or modems that communicate with external time sources like the Rugby Atomic Clock. You *must* have at least one Primary time server if you have a Reference server.	Hardware clock or external source.	No	Primary, Secondary, and clients
Secondary	No	Default configuration. Attempts to remain synchronized with only one time source. Does not participate in voting.	Single Reference, Reference, or Primary.	Yes (100% correction per polling interval)	Clients only

Table 4-2: Time Server Types

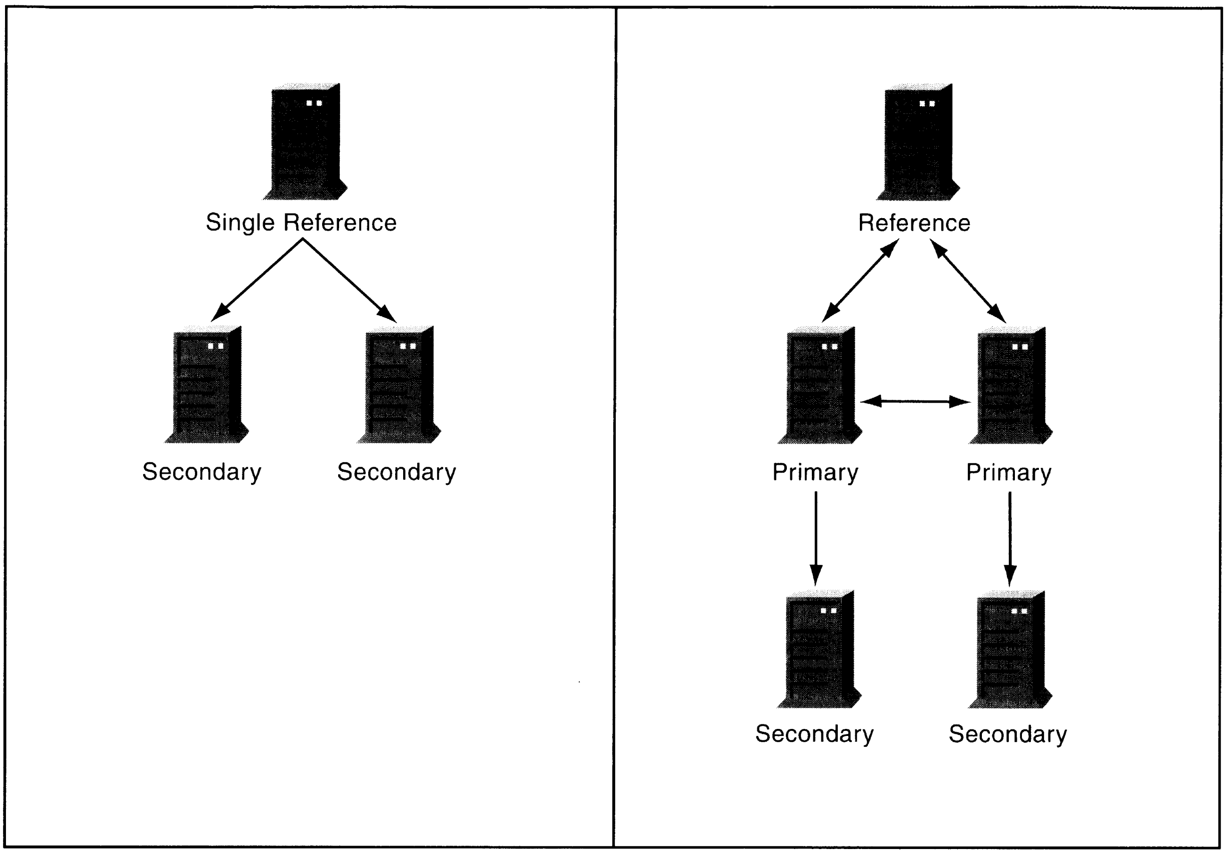

Figure 4-11

Two Methods for Time Synchronization

You can use two methods to keep NetWare 4 servers synchronized:

- Default configuration

  This method uses SAP (Service Advertising Protocol) to transmit and receive time.

- Custom configuration

If you have a network with a small number of servers, use the default configuration for time synchronization services. If you are using an organizational implementation, you may want to use a custom configuration to decrease network traffic.

The subsections below explain these methods in detail, including the advantages and disadvantages of each.

For more information on SAP, refer to *Supervising the Network.*

Default Method

The installation program for NetWare 4 assumes that only two types of time servers are necessary: Single Reference servers and Secondary servers.

The default installation method was provided for two reasons:

- It is simple and efficient.
- It does not require you to reconfigure time synchronization when new servers are added to the network.

The default method uses the Service Advertising Protocol (SAP) to advertise the time sources (Primary, Reference, and Single Reference servers). Using SAP, Reference and Primary time servers know which servers to arbitrate time with, and Secondary servers know where to get their time from without intervention from the network administrator.

The default method may not be ideal for organizational implementations with many sites connected by WAN links. If each site has two or more Primary servers, the voting process may involve more network traffic than necessary.

Advantages

The default method has the following advantages:

- It is easy to understand and requires no planning.
- It does not require a configuration file.

  Because the default configuration relies on SAP to advertise the time source, no configuration information needs to be provided to any of the time servers.
- The chance of synchronization error is reduced because a time receiver (Secondary server) will only talk to a time provider (Reference, Single Reference, or Primary) and never to another Secondary for its time.

Disadvantages

The default method has the following disadvantages:

- The Single Reference server must be contacted by every other server on the network.
- Using SAP means that a misconfigured server can possibly disrupt the network. Some of the Secondary servers may synchronize to an unsynchronized server rather than to the authorized Single Reference server.
- One time source means a single point of failure. However, should a Single Reference server go down, a Secondary time server can easily be set as the Single Reference server using a SET parameter or using SERVMAN.NLM.

Custom Configuration

Custom configuration is not difficult, but it does require some planning. It uses Reference and Primary servers as time providers to minimize a single point of failure.

To use a custom configuration, you must do the following:

- Determine which servers are time sources.
- Decide which time source individual servers will refer to.

Each server is then given a configuration file (TIMESYNC.CFG) listing the authorized time sources for the server and other parameters. The same configuration file can often be copied to several servers. Often the only change required is the order of the time sources.

The custom configuration cuts down on network traffic, but it requires additional administration. Therefore, we do not recommend it for departmental implementations.

Advantages

The custom configuration has the following advantages:

- You have complete control of the time synchronization hierarchy.
- You can optimize network traffic and distribute time sources around the network.
- You can provide alternate time sources to be used in case of network failures.

Disadvantages

The custom configuration has the following disadvantages:

- Customization requires careful planning, especially on a large network.
- Adding new time sources usually requires that configuration files on several servers be updated.

SET Parameters for Time Synchronization

Time synchronization has two SET parameters that allow you to further customize your time services.

Directory Tree Mode

When the Directory Tree mode is set to ON, time servers throughout the Directory tree will only listen to and vote with time servers in their own Directory tree (if other trees exist on the same network).

SAP Mode

When the SAP mode is set to ON, time servers will receive and send time broadcasts over SAP. You can use this with mode configured lists; if a server cannot contact any server in its list, the server will listen to SAP for its time.

Directory Tree mode and SAP mode can be used together to increase the likelihood that time servers will always find another time server to vote with or poll. Using both of these modes can also help reduce the administration of the configured lists.

Guidelines for Using a Custom Configuration

When customizing time synchronization, the general approach is to create a hierarchical structure based on the physical location of servers. Use the following guidelines:

- Have Secondary servers synchronize to time sources—Primary, Reference, or Single Reference servers.
- Keep the number of time sources as small as possible to reduce network traffic. These sources should be high-visibility servers, rather than satellites, on the network. You should have no more than five Primary and Reference servers.
- Use time sources to provide local access throughout the network.

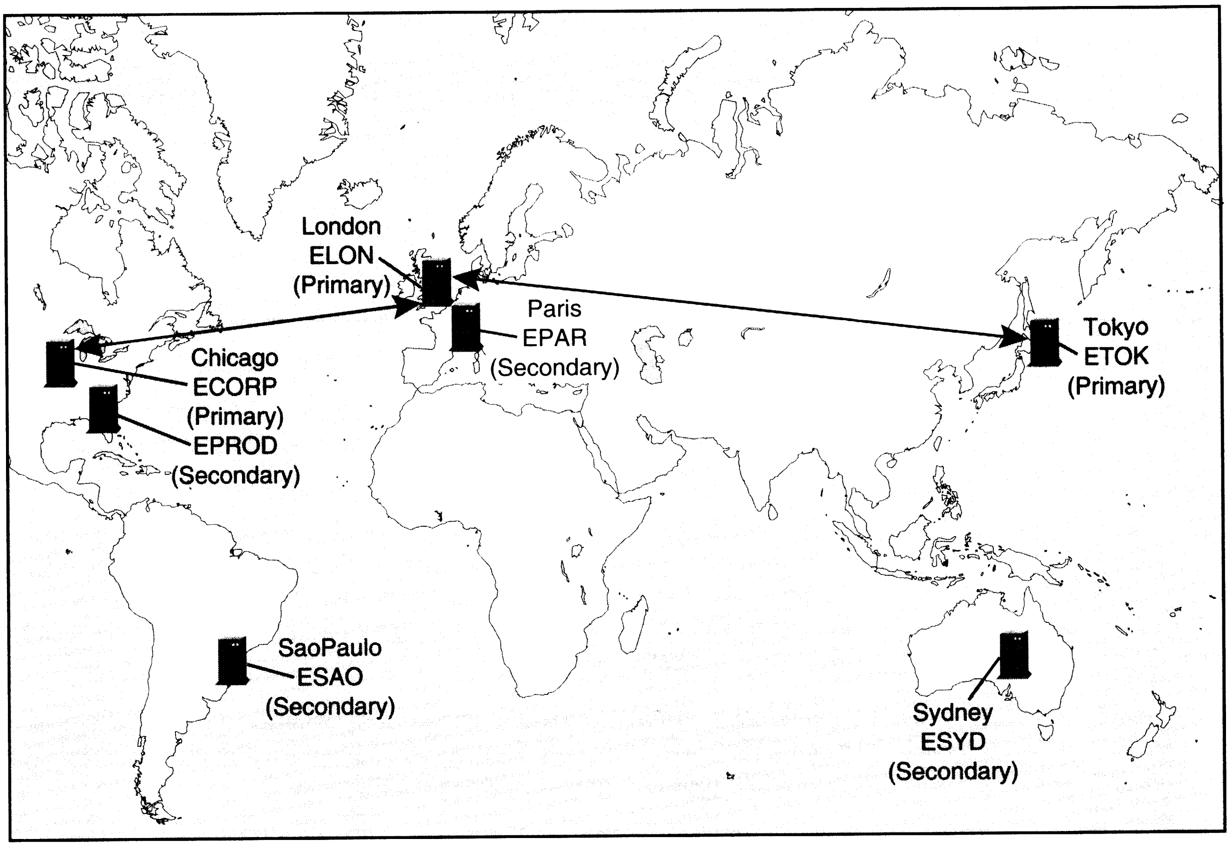

Figure 4-12

EMA Time Synchronization

The case company, EMA, can be used to illustrate a custom time synchronization configuration. The EMA time synchronization setup includes a small set of time-providing servers, configured to enhance fault tolerance and performance of time services. The time providers consist of Primary servers in each geographical region that are responsible for time synchronization across all the sites. Other servers look to these regional Primary servers as the source of their time.

For example, EMA wants to coordinate time synchronization services among the currently installed locations. They decide on the following:

Chicago/ECORP	Reference
Tokyo/ETOK	Primary
London/ELON	Primary
Others	Secondary

Table 4-3: EMA Time Synchronization Setup

Time is arbitrated among the Primary servers and the Reference server to agree on a correct time. Secondary servers get their time from the closest time source.

In each geographic region, the time source list should be ordered so that the closest time provider appears first in the list. The least-cost providers (in terms of communication link speed) appear next, and the rest of the providers follow. This allows the Secondary server to obtain the time from the closest available server.

If you know that the Secondary servers that communicate with the time sources are good time keepers, you could increase the polling interval to an hour or more to reduce traffic across the links.

In the following example, we will increase the ESAO server's polling interval to one hour.

Sample TIMESYNC.CFG Files

In this example, the TIMESYNC.CFG file for Reference server ECORP contains the following parameters:

```
#Configuration Parameters from server ECORP

Configured Sources = OFF
Directory Tree Mode = ON
Hardware Clock = ON
Polling Count = 3
Polling Interval = 600
Service Advertising = ON
Synchronization Radius = 2000
Type = REFERENCE

#Configured time source list from server ECORP
```

The TIMESYNC.CFG files for the other servers contain the following parameters:

Configured Sources	ON	ON	ON	ON	ON	ON
Directory Tree Mode	ON	ON	ON	ON	ON	ON
Hardware Clock	OFF	OFF	OFF	OFF	OFF	OFF
Polling Count	3	3	3	3	3	3
Polling Interval	600	600	600	600	600	600
Service Advertising	ON	ON	ON	ON	ON	ON
Synchronization Radius	2000	2000	2000	2000	2000	2000
Type	Secondary	Primary	Secondary	Secondary	Secondary	Primary
Time Source	ECORP	ECORP ETOK	ELON	ECORP	ETOK	ECORP ELON

Table 4-4: Sample TIMESYNC.CFG Parameters for EMA

Managing Time Synchronization

This subsection discusses additional information you need to know to set and manage time synchronization on your network.

Changing Time

If a network contains only Primary servers as time sources, setting the network time is very difficult. Setting the time on a Primary server causes the server to believe it is not synchronized, and it tries to correct the error. Time synchronization may fight efforts to correct misconfigured local parameters.

To avoid this problem, use the following command after setting the server time and creating a TIMESYNC.CFG file:

```
SET TIMESYNC RESTART FLAG = ON
```

This should be done on only one of the Primary servers. The hardware clock then becomes like an "external time source" for the network.

You can also configure a Reference server to get its time from its own hardware clock. This is an inexpensive and reliable way, for the most part, to synchronize time to an external source.

Setting hardware time is strongly discouraged for most applications that require synchronized time stamps. To make sure that the hardware clock has the correct time, always set the DOS time before booting the server.

If a time server's clock is set back, other servers will not set their clocks back. Instead, they will advance their clocks at a slower pace and the slower server will advance its clock at a faster pace until the servers converge.

Do not reset server time carelessly. If you accidentally reset time to a future year, for instance, you risk serious corruption of the Directory from erroneous time stamps.

Keep in mind that the reason for time synchronization is to keep the network clock stable by having the servers report the same time. Consensus among time servers is more critical to NDS functions than is the reporting of true and accurate time. However, other services that rely on time will suffer if the clock is incorrect.

Creating Multiple Time Provider Groups

For networks with a large number of servers at more than one location, you may want to consider creating a time provider group at each location. The benefit of such a configuration is reduced time synchronization traffic across WAN links.

A time provider group should include a Reference server and a small number of Primary servers.

Reference Time Servers

For instance, if a company had 12 servers in London and 23 servers in Frankfurt, you might have one Reference server at each location using a dial-up time service. Thus, the Reference servers would synchronize to a common clock without generating any synchronization traffic across the WAN.

Primary Time Servers

A few Primary servers could then be used in London and a similar number in Frankfurt. Each Primary time server would use the local Reference server as its time source.

Secondary Time Servers

All other servers would be Secondary time servers. Each Secondary time server would list the local Primary time servers as time sources in its TIMESYNC.CFG file.

Recommendations for Bringing a NetWare 4 Server Online

When you bring a NetWare 4 server online, one of the first things the server checks for is the network time. Because the server assumes that the network time is the correct time and that its own clock is not synchronized with the network time, the server sets its clock to the network time.

This ensures that you will not disrupt time services on the network by bringing up a misconfigured time source. All time providers will adjust their clocks 100 percent when they come online with the network, with a few small exceptions.

You can minimize time synchronization problems by always performing the following steps when you bring a server online:

1. Before launching SERVER.EXE, set the DOS (hardware) clock to the appropriate local time.
2. When the server console comes up, check the server's Time Synchronization information to make sure that the time zone and daylight savings values found in the AUTOEXEC.NCF file are correct.
3. Observe the server for a few minutes, checking its time from the console two or three times to make sure it does not adjust by an hour or more.

If the server's local console time does adjust by more than an hour during the few minutes you are observing, you might assume that you have incorrectly set one of the values for UTC. Take immediate action to correct the values, and watch the server's time again. Do not let the server run online for more than 20 minutes with a drastically incorrect time, or you may cause unpredictable problems with the network clock.

Exercise 4-2: Planning Time Synchronization for Departmental Implementation

Scenario: The Human Resources department at EMA is implementing a NetWare 4 network. The department will implement its own tree until IS is ready to incorporate the department into the corporate tree.

Procedure: Given the preceding scenario and based on the following servers being installed, plan Human Resource's time synchronization strategy:

[illegible]	Order of Installation
HR-ADMIN	First installation
HR-PERSONNEL	Second installation
HR-RECRUITING	Third installation

1. What will be the default time server types?

[illegible]	Type of Time Server
HR-ADMIN	
HR-PERSONNEL	
HR-RECRUITING	

2. What type of synchronization method (default or custom) would you implement and why?

3. When might you consider installing a Primary server instead of a Single Reference server?

Exercise 4-3: Changing Time Synchronization

Scenario: The following table shows the time sources for EMA.

Location/Server	Type of Time Server
New York/ ECORP	Reference
London/ELON	Primary
Tokyo/ETOK	Primary

These servers were selected because of their reliable clocks and centralized locations on the EMA network topology.

Procedure: Given the preceding scenario and based on the following table, implement the time synchronization strategy.

	Server						
Parameter	**ECORP**	**EPROD**	**ELON**	**EPAR**	**ESAO**	**ESYD**	**ETOK**
Configured Sources	ON	ON	ON	ON	ON	ON	ON
Directory Tree Mode	ON	ON	ON	ON	ON	ON	ON
Hardware Clock	ON	OFF	OFF	OFF	OFF	OFF	OFF
Polling Count	3*	3*	3*	3*	3*	3*	3*
Polling Interval	10*	10*	10*	10*	60	10*	10*
Service Advertising	OFF	ON	ON	ON	ON	ON	ON
Synch. Radius	2000*	2000*	2000*	2000*	2000*	2000*	2000*
Type	Reference	Secondary	Primary	Secondary	Secondary	Secondary	Primary
Time Source	ETOK ELON	ECORP	ECORP ETOK	ELON	ECORP	ETOK	ECORP ELON

* If the server is communicating over an encrypted WAN link, or if it "times out", or if you experience synchronization failure errors, you can increase these values by three-fold or more.

1. At your server console, launch the SERVMAN utility. Type

   LOAD SERVMAN <Enter>

2. Select **Server parameters**.
3. Select **Time**.
4. Highlight **TIMESYNC RESET** and press <Enter>.
5. Select **Yes**.
6. Change the parameters for your server using the values from the table on the previous page.
7. Highlight **TIMESYNC Write Parameters** and press <Enter>.
8. Select **Yes**.
9. Highlight **TIMESYNC Restart Flag** and press <Enter>.
10. Select **Yes**.
11. Press **<Esc>** to exit SERVMAN; answer **Yes** to update the TIMESYNC.CFG file now.
12. To view the options you have set, type the following at the server console:

    LOAD EDIT TIMESYNC.CFG <Enter>

13. If you need to change the TIMESYNC.CFG file, do the following:

    a. Edit the file.

    b. Press **<Esc>** and confirm **Yes** to save the file.

    c. Press **<Esc>** again to exit the editor.

    d. To use the new parameters in TIMESYNC.CFG, type

       SET TIMESYNC RESTART FLAG = ON <Enter>

Summary

NDS partitioning and replicating provides protection for and enhances performance of NDS.

Time synchronization is critical in maintaining database integrity because it ensures that NDS operations occur in the proper sequence.

SECTION 5 Creating a Detailed Design and Troubleshooting NDS

Introduction

In this section, you will learn how to create a detailed Directory tree design. You will also learn how to recover from hardware failures, which can adversely affect NDS. You will learn how to prepare the server for downtime and remove it from the Directory tree. These procedures will help you properly repair the Directory database in the event of data corruption or loss.

Objectives

At the end of this section, you will be able to do the following:

1. Create a detailed design of the Directory tree that provides for securing NDS, partitioning the Directory, replicating partitions of the Directory, and synchronizing time.
2. Identify NDS inconsistencies.
3. Execute the procedures to prepare an NDS server for downtime.
4. Execute the procedures to recover from a failed master replica.
5. Perform the steps to remove a server from the Directory tree.

Detailed Design Issues

In Section 2, "Designing and Administering NDS," you planned a structural design of a Directory tree. This section will focus on the detailed design of the Directory tree. You will determine how security, replication, and synchronization affect the Directory tree structure (see Figure 5-1).

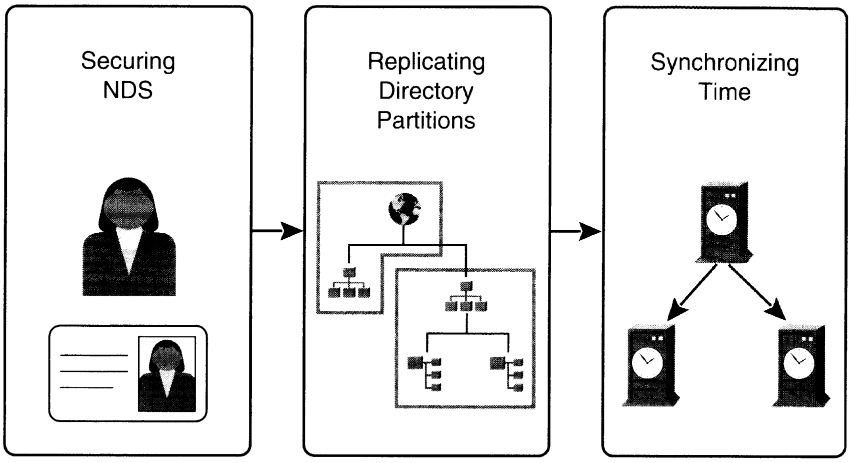

Figure 5-1: Detailed Design

Before reviewing the structural design and applying the detailed design decisions, make sure you understand the information and procedures covered in Section 3, "Securing the Directory Tree," and Section 4, "Partitioning and Replicating NDS and Synchronizing Time."

Table 5-1, on the following page, summarizes the concepts, procedures, defaults, and additional needs for each of these areas.

	Securing NDS (Section 3)	Partitioning and Replicating (Section 4)	Synchronizing Time (Section 5)
Concepts and Procedures	❑ Rights assignment ❑ Access restrictions ❑ Container administrators ❑ User rights ❑ Profile ❑ Directory Map ❑ Alias ❑ Security equivalence	❑ Partition ■ [Root] partition ■ Partition root or master ❑ Replicas ■ Master ■ Read/write ■ Read-only ■ Subordinate reference ❑ Replica ring ❑ Moving containers ❑ Parent and child replicas	❑ Time servers ■ Time Providers • Single Reference • Reference • Primary ■ Secondary (time receiver) ❑ Time Provider group ❑ Time source ❑ Universal Time Coordinated (UTC) ❑ TIMESYNC.CFG
Defaults	❑ Admin ❑ Public ❑ Server installation ❑ User creation ❑ Container rights	❑ Same partition ■ Installing first server ■ Installing second server ■ Installing third server ■ Installing fourth server ❑ Subordinate reference ❑ Merge servers	❑ First server vs. additional servers ■ Installation ■ Upgrade ■ Merge ❑ Single Reference
Additional Needs	❑ Determine centralized or distributed administration ❑ Provide rights for ■ Profile login scripts ■ Directory Map ■ Alias objects ■ Group access to containers ■ Limited access within containers ■ Users in different context ■ Traveling users	❑ Splitting partitions ❑ Joining partitions ❑ Moving partitions ❑ Changing replica type ■ Assigning read/write replicas ■ Using read-only replicas ■ Changing master replica location ❑ Fault tolerance ■ Replica of [Root] ❑ Provide quick user access	❑ Merging trees ■ Source ■ Type ■ Write parameters ■ Restart flag ❑ Providing different time source ❑ Creating time provider groups ❑ Changing communication method (Directory and SAP) ❑ Changing server type

Table 5-1: Summarizing Security, Partitioning, and Synchronization

After reviewing the information in Table 5-1, consider the design issues for the Directory tree. Table 5-2 summarizes several design issues.

	Securing NDS	Partitioning and Replicating	Synchronizing Time
Detailed Design Issues	❑ Determine administration approach: distributed vs. centralized ❑ Place groups in the Directory ❑ Plan inheritance and security equivalences ❑ Assign Server object rights ❑ Provide traveling user access	❑ Plan partition boundaries ❑ Identify the appropriate replica assignments ❑ Determine accessibility and fault tolerance vs. network traffic and performance ❑ Allow for WAN links ❑ Assign administrators of partitions and replicas	❑ Plan time server types for your Directory tree ❑ Coordinate time source servers and time provider servers ❑ Reduce WAN traffic for synchronizing time ❑ Plan a time synchronization strategy for merges

Table 5-2: Summarizing Detailed Design Issues

These issues will be explained, in detail, in the next several pages.

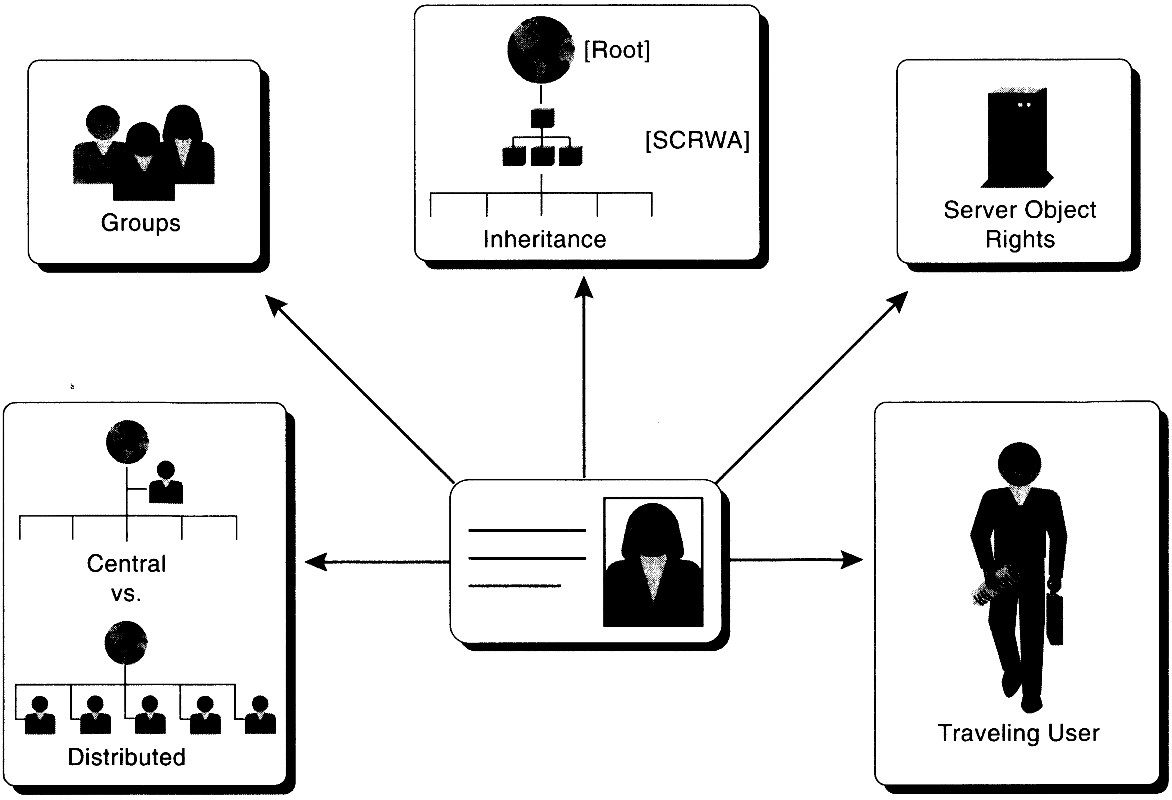

Figure 5-2

Securing NDS

You should consider several design issues when you determine how to secure the Directory tree.

Determining the Administration Approach

You can use the centralized or the distributed approach to administering security. If you use the centralized approach, the default setup works with the Admin object that is created with the first server. If you use the distributed approach, you will need to set up container administrators with the appropriate rights after you have created the structure and installed the remaining servers.

For further information, refer to “Developing a Security Strategy” in *Introduction to NetWare Directory Services.*

Placing Groups in the Directory

You should determine whether to provide containers for workgroups that are small or use Group objects to provide access to resources. If you have a small company, you can use containers to provide access to the resources. If you have a larger company and several workgroups are using the server, you can use Group objects to provide the resources to specific people.

If members of a group are spread across a WAN, members will authenticate across the WAN with each login.

Planning Inheritance and Security Equivalences

Your group and container design affects what rights you allow objects to inherit and which objects require security equivalence. The rights that you want to flow down the tree are critical to your Directory tree design.

Also, each container above a User object (toward the [Root]) supplies rights to the object. These rights from the container are known as implied security equivalences but are not shown in the Security Equal To property of the User object.

Assigning Server Object Rights

When you assign rights to the Server object, those rights flow down to the file system. Be careful not to assign a container, Group, Organizational Role, or any other object the Supervisor object right. The Supervisor right will flow to the file system.

Example: A network administrator was wondering why all the User objects have supervisor access to the file system.The administrator determined that the source of these rights came from assigning the container object Supervisor rights to the Server object.

Providing Traveling User Access

Several approaches can be used to provide access to network resources for the traveling user. Make sure you design your tree with these users in mind. Refer to the different approaches in Section 3, "Securing the Directory Tree."

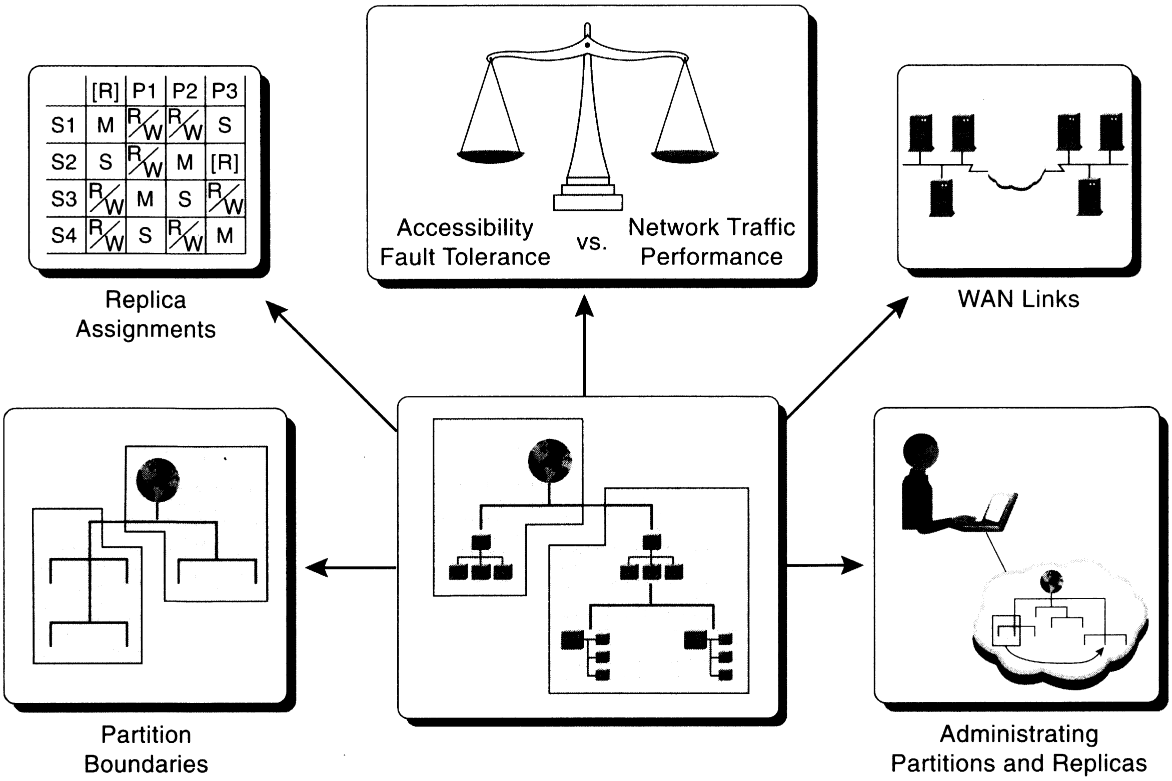

	[R]	P1	P2	P3
S1	M	R/W	R/W	S
S2	S	R/W	M	[R]
S3	R/W	M	S	R/W
S4	R/W	S	R/W	M

Figure 5-3

Partitioning and Replicating

You should consider several design issues when you determine partitions and replicate them in the Directory tree.

Plan Partition Boundaries

From your logical Directory tree structure, determine where your partitions will exist. The critical factors when making the decision to place partition boundaries include WAN topology, geographic location, access of Directory information, number of objects in the containers, workgroup needs, information flows, elimination of single points of failure, and reduction of unnecessary network traffic.

For further information, refer to “Developing a Replication Strategy” in *Introduction to NetWare Directory Services*.

Identify the Appropriate Replica Assignments

The default replica assignments help prevent single points of failure; however, you may want to plan extra assignments for accessibility. Use the guidelines in Section 4, "Partitioning and Replicating NDS and Synchronizing Time," to determine the number of replicas, the types of replicas, and the servers needed for storing the replicas to maximize fault tolerance and accessibility.

Determine Accessibility and Fault Tolerance vs. Network Traffic and Performance

The main concerns for partitioning and replicating are to have enough replicas of your Directory tree to make the data available when needed and to provide fault tolerance in case a master replica is corrupted. An optimal goal is to provide fault tolerance and accessibility without decreasing performance by requiring too much network traffic to synchronize the replicas.

Allow for WAN Links

Too much network traffic is especially critical for WAN links. Make sure you minimize the replica synchronization across expensive lines. Consider making smaller partitions on fewer servers to reduce traffic.

Assign Administrators of Partitions and Replicas

Partitioning changes create network traffic, especially if several actions are requested at the same time. To avoid this, determine who will make partition changes. Make sure any changes are coordinated to prevent excessive network traffic.

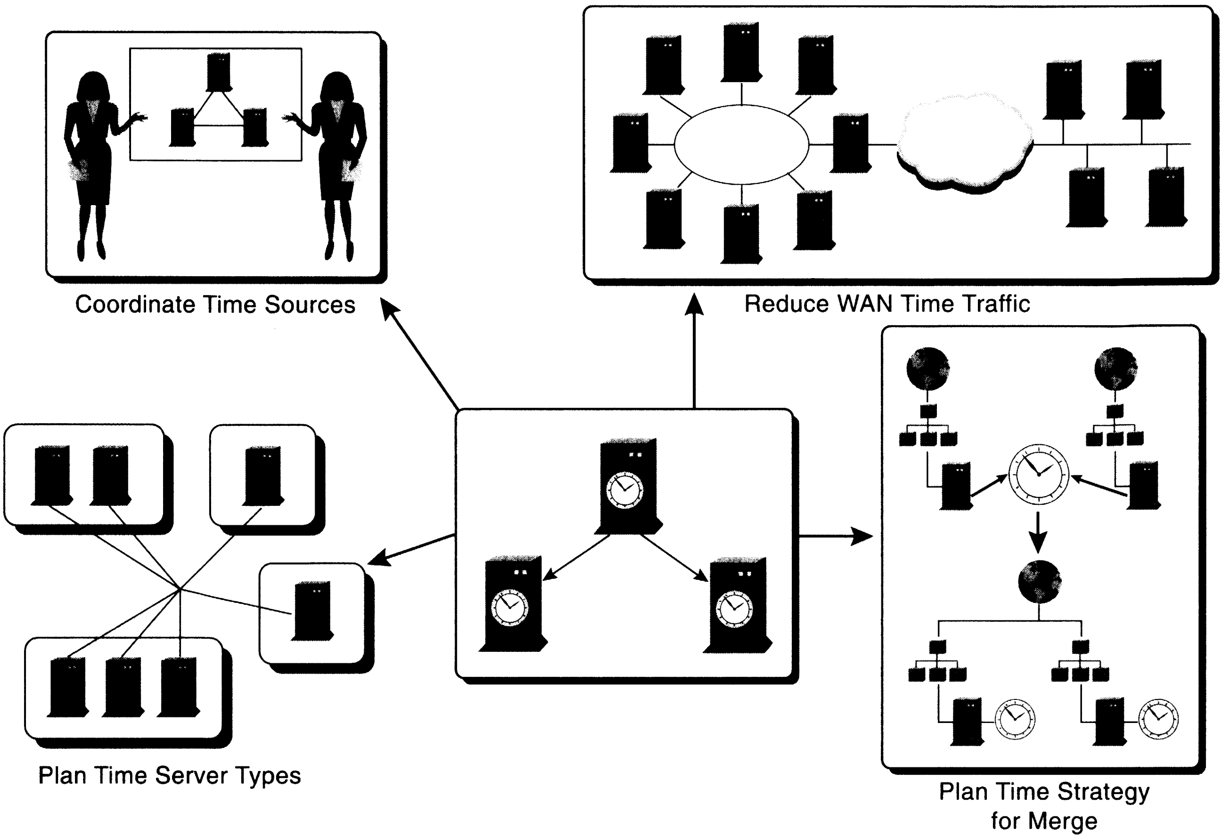

Figure 5-4

Synchronizing Time

You need to consider several design issues when assigning time synchronization.

Plan Time Server Types for Your Directory Tree

You need to determine the types of time servers that are most appropriate for your situation. For small workgroups, the default settings will suffice: The first server is a Single Reference server and the subsequent servers are Secondary. When you have a divisional or organizational Directory tree, you should plan to use both time provider servers and time receiver (secondary) servers. Keep in mind the possible conflicts and network traffic requirements.

For further information, refer to “Developing a Time Synchronization Strategy” in *Introduction to NetWare Directory Services.*

Provide the plan for all network administrators if you are installing an organization-wide Directory tree. This plan will help coordinate and determine time server types when servers are added to the Directory tree.

Coordinate Time Sources

When you have a divisional or organizational Directory tree, you need to determine the source servers or time provider groups. Review the advantages of the server types to maximize synchronization and minimize network traffic. Make sure the plan for time server type is coordinated with other container administrators if your NDS administration is distributed.

Reduce WAN Traffic for Synchronizing Time

If you have a large Directory tree that spans WAN links, make sure that the time providers do not have to communicate with secondary servers across the WAN link.

Plan a Time Synchronization Strategy for Merges

Prior to the merge, you are required to have your servers synchronized. Determine an appropriate plan for time synchronization before the merge. To avoid conflicts, refer to Section 4, "Partitioning and Replicating NDS and Synchronizing Time," for information about time server types and their characteristics. For example, do not have more that one Single Reference server on the same Directory tree.

Written Exercise 5-1: Create a Detailed NDS Design

Activity: Implement the factors in your detailed design of your Directory tree.

Procedure: Using the scenario assigned by your instructor, implement a design for a Directory tree that includes these components: securing NDS, partitioning the Directory, replicating partitions of the Directory, and synchronizing time across the Directory.

1. You will work in one of three groups to review the structural design and determine implications from the components for a detailed design.
2. Each group will work on the same scenario for which they created a structural design.
3. Each group will present the information to the class and will include rationale for decisions. The information in the presentation should include the following:
    - Security strategy
    - Administration strategy
    - Partition boundaries
    - Replica assignments
    - Time providers
    - Server time type

To prepare your presentation, refer to examples of small, medium, and large Directory trees in *Introduction to NetWare Directory Services.* These examples include the types of information that you should present to the class.

Scenario 1: Using the Directory tree for TCabinets (from Exercise 2-2), determine the partition boundaries, replica assignments, time server types, and security strategy.

Servers/Type	**Partitions**			
	[Root]			
Reference	Master			

Draw the Directory tree:

Scenario 2: Using the Directory tree for Standard Publishing (from Exercise 2-2), determine the partition boundaries, replica assignments, time server types, and security strategy.

Servers	**Time Type**	**Partitions**		
		[Root]		
	Single	Master		

Draw the Directory tree:

Scenario 3: Use the Directory tree for Sports, Inc. (from Exercise 2-2). Sports, Inc. has added two more servers. Provide replica placement and time server types for the two new servers. Determine the partition boundaries and security strategy.

Servers	**Time Type**	**Partitions**		
		[Root]		
		Master		

Draw the Directory tree, showing the additions:

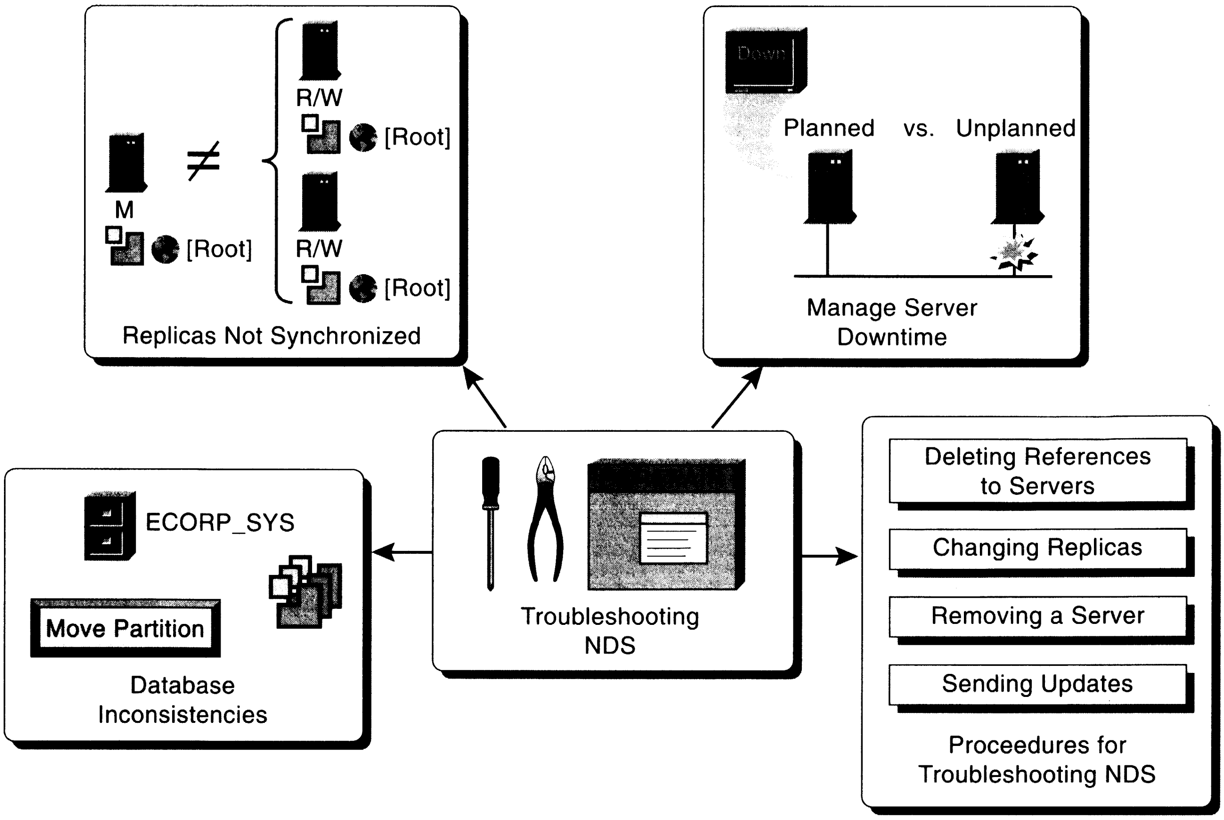

Figure 5-5

Troubleshooting NDS

Now that you have learned the NDS components, you can determine how to deal with problems you may experience using the Directory. NDS provides tools and strategies for troubleshooting the following areas:

- Directory database inconsistencies
- Unsynchronized replicas
- Server downtime

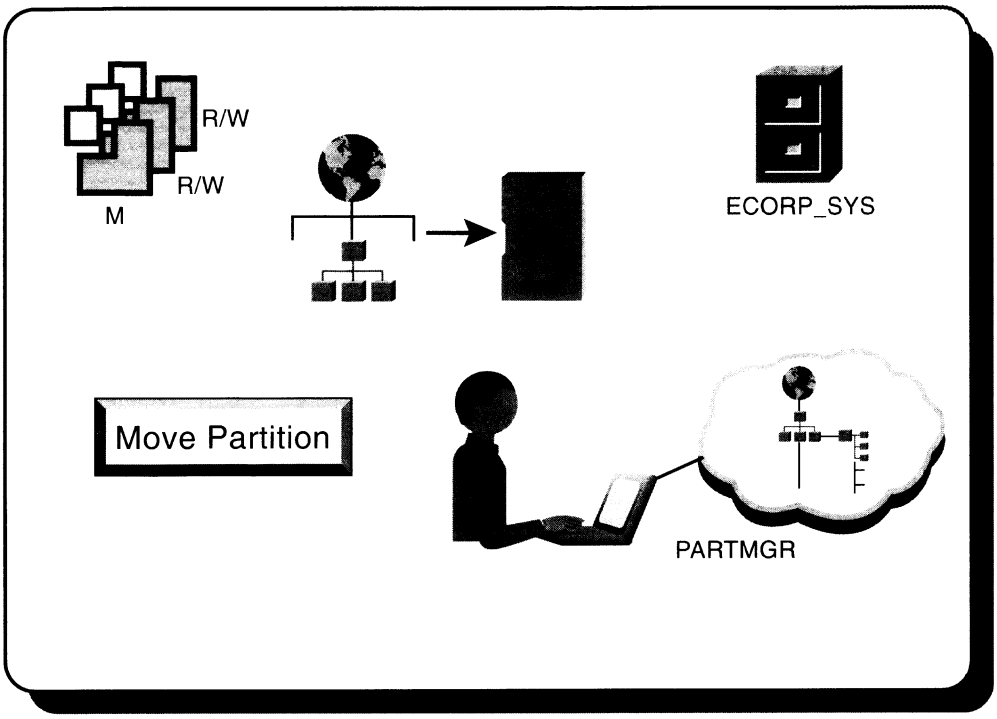

Figure 5-6

Managing Directory Database Inconsistencies

NDS is a distributed, replicated database that is loosely consistent. The database requires time to replicate and synchronize major network changes.

The following subsections explain how the database works and suggest how you can prevent inconsistencies in the database.

Understand Replication and Synchronization of Changes

The amount of time required for a change to be replicated and synchronized depends on the type of change, the size of the partition, and the number of servers that the partition is replicated on.

You should not assume that delays in replication and synchronization of changes indicate database inconsistencies.

Simple Changes

A change that affects only a single User object (such as changing the phone number) takes relatively little time to be replicated and synchronized, since the change is sent only to the file servers that contain a replica that includes this User object.

Creating a partition takes very little time as well. When you create a partition, the Directory tree uses partition attributes to "draw" the new partition boundary. The information needed is already replicated around the network.

Complex Changes

More complex changes take more time. For example, joining two partitions that exist on two different sets of servers will take time. Each server that has a piece of the original partitions may receive a replica of each partition before the join is performed. In this example, NDS performs the following tasks:

1. NDS automatically determines where all the replicas of each partition exist.
2. NDS forces the servers to replicate the data of both partitions to the servers.
3. The join is completed; the servers have the composite information of both partitions.

Ensure Partition Fault Tolerance

To ensure partition fault tolerance, you should always have three replicas of every partition in the tree. These replicas provide the best backup of the data stored in the Directory database. Make sure you have replicas located on-site and off-site; make sure regular backups are stored on- and off-site as well.

Back Up the Directory

You should always back up the Directory before you split or merge a partition. You should back up the Directory on a regular (weekly or monthly) basis and keep the backup as a second protection for restoring the Directory.

If some critical interruption (such as a power failure) occurs during the merge or split, the Master replica for a partition may become inconsistent with other replicas. It can then spread errors to other servers containing replicas of that partition; these servers will accept the changes as the latest changes.

The first recovery option in such a case would be to make a read/write replica from an off-site server into the master replica. The next option would be to use your Directory backup.

Manage Partitions from One Workstation

You should perform partition merges and splits from only one workstation at a time. This helps you track the partitions.

In NetWare 4.1, a partition locks during a merge, move, or join; therefore, two simultaneous operations are impossible. Generally, the merge or split locks the partition for a few seconds to a few minutes.

Do Not Let the SYS: Volume Run Out of Space

The SYS: volume contains the NDS database in a hidden file. This directory is not viewable even with the NDIR command. The NDS database is protected by the Transaction Tracking System (TTS). If the SYS: volume becomes full, TTS shuts down and the Directory closes its files to any changes.

To avoid running out of disk space on the SYS: volume, use the following guidelines:

- Set minimum space requirements so that you will receive a warning if the SYS: volume is almost out of space.
- Store print queues on a volume other than the SYS: volume.
- Store user files or other files on a volume other than the SYS: volume.
- Do not add replicas to servers that are low on disk space.
- If you have CD-ROM drives attached to the server, they often create large index files on the SYS: volume. Make sure the volume is large enough to accommodate these index files.
- If auditing is enabled, make sure there is an acceptable limit on the auditing file size.
- Use the SET AUTO TTS BACKOUT FLAG = ON in the STARTUP.NCF file. This flag backs out any incomplete TTS transactions to prevent inconsistencies in the database.

TTS must always be running on a server that contains active replicas. If TTS shuts down, unpredictable and possibly damaging results may occur to the replicas on that server.

Determining If Replicas Are Not Synchronized

It is important to remember that the Directory database is a loosely consistent database and it takes time to replicate and synchronize changes. However, persistent problems may indicate that replicas are out of synchronization.

Client Symptoms

The following client problems may indicate that replicas are out of synchronization:

- The client prompts for a password when none exists for the user account.
- Modifications to the Directory seem to have "disappeared."
- NDS rights assigned previously seem to have "disappeared."
- Client performance is inconsistent; errors cannot always be duplicated.

Server Symptoms

You can use a SET parameter named DSTRACE to watch the synchronization of the Directory database. Enter SET DSTRACE=ON at the server console. This brings up another screen on which you can watch the replicas synchronize with each other.

If you get the following message on a regular basis and it persists for more than 20 minutes, you should take the actions listed below to correct the problem:

SYNC: End sync of partition <*name*>. All processed = NO

Corrective actions:

1. Let the system run for a few hours. It may synchronize and correct itself. *Do not bring down the file server; you will prevent self-correction.*
2. Run DSREPAIR.NLM. The steps for further determining the state of partitions are covered in "Procedures for Troubleshooting NDS," in this section.

Do not attempt to split or join a partition that is already experiencing problems. You may cause more synchronization problems.

Unknown Object Types

The presence of Unknown objects in the Directory tree can indicate a problem with synchronization. However, Unknown objects do not always indicate a problem.

Sometimes objects become unknown during the Merge and Create Partition operations. This is normal because a partition root is changing. The objects will return to a known state when the operation is completed.

Managing Server Downtime

This subsection discusses how server downtime affects replicas and explains how you can maintain database integrity. Specific considerations are different for the following areas:

- Downtime due to planned maintenance
- Downtime due to failure

Down Time Due to Planned Maintenance

The Directory database is designed to withstand periods during which servers that store replicas are unavailable. The Directory on the unavailable server automatically resynchronizes itself when it becomes available again.

Servers that have changes to report to the unavailable server have a backoff algorithm that tells them not to send changes because the server is not connected.

The disconnected server also has a routine that periodically requests changes if no changes have been received for a set period of time. The downed server relies on this algorithm to synchronize its database once it comes back online.

If a downed server contains the master replica, the services and user accounts on that partition are not available until the server is reconnected. If you plan to bring down a server that contains the only replica of a partition, move that partition to another server using the Partition Manager utility.

Event	Action
Bringing down a server or WAN link permanently.	Back up NDS; then remove NDS from the server using INSTALL.NLM.
Bringing down a server or WAN link temporarily to replace the hard drive that contains the SYS: volume.	
Bringing down a server or WAN link temporarily if any of the following are required: ■ Partition operations ■ Large numbers of changes, additions, or deletions of replicas ■ Relocation of the server to another site	Remove replicas from the server before bringing down the server or disconnecting WAN links. (Bindery services is lost when replicas are removed.)

Table 5-3: Managing NDS During the Planned Shutdown of a Server or WAN Link

Down Time Due to Failure

A server that fails (because of hard-disk loss or some other problem) can cause Directory loss if that server contains replicas of the Directory database.

Normally, users should be able to access network resources because you have other replicas of the partitions stored on other servers.

However, if a server fails, you should remove the server from the Directory tree using Partition Manager. This action will let the other servers know immediately that the server is no longer part of the replica ring. Otherwise, servers with the same replicas may try indefinitely to update the failed server's replicas.

Eventually, the other servers in the ring send updates less frequently, but the servers will continue to attempt to contact the server until it is deleted or comes back online.

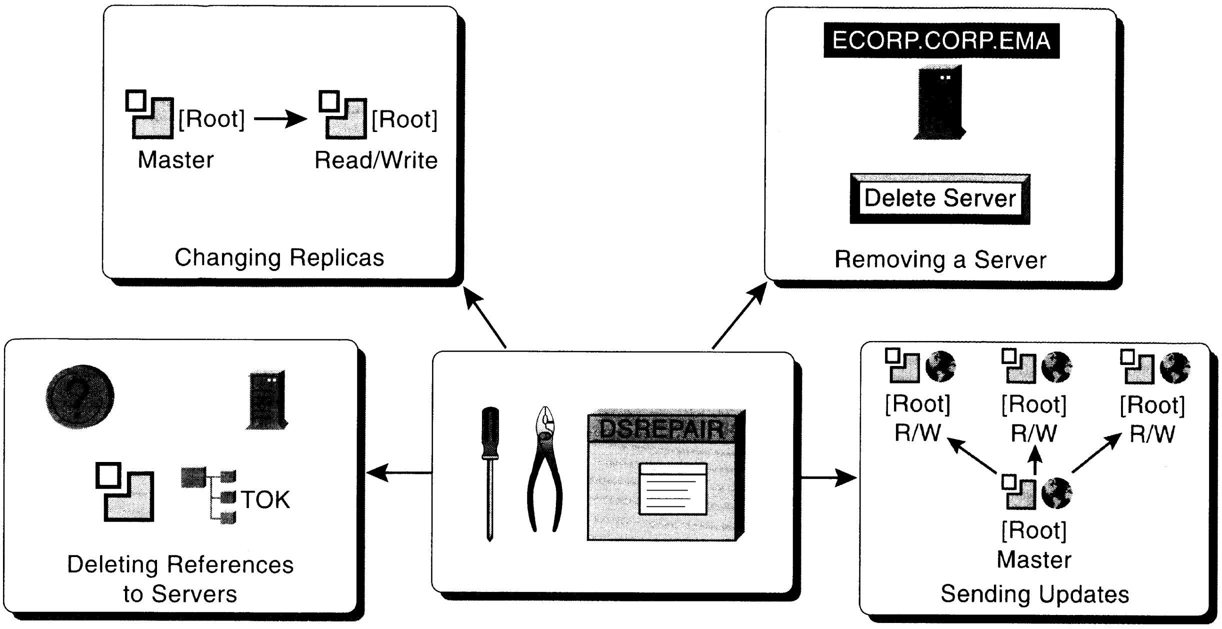

Figure 5-7

Procedures for Troubleshooting NDS

You can use the following procedures to troubleshoot NDS:

- Determine the synchronization state of a partition.
- Restore system software and NDS after a SYS: volume failure.
- Remove a server from the Directory tree.
- Send updates to synchronize corrupted replicas.

Determining the Synchronization State of a Partition

Because of the global nature of NDS, you should check a partition's synchronization state before performing certain NDS operations. To determine a partition's synchronization state, follow these steps at the console:

1. Type

   LOAD DSREPAIR <Enter>

2. Select **Replica Synchronization**.
3. Enter the administrator's full distinguished name and password.
4. Check the DSREPAIR log file for the status (OK) of each server.

   You can perform NDS operations safely when you receive this message for all the partitions on that server.

Restoring System Software and NDS after a SYS: Volume Failure

Because the Directory is located on the SYS: volume, a hard drive crash involving the SYS: volume is equivalent to removing NetWare 4 from the file server. You will need to reinstall NetWare 4 and NDS before you restore your data.

If the server that went down contained the master replica for a partition, there are further complications. The master replica must be available during partitioning changes. If the server hard disk that contains a master replica crashes, you must find another replica and upgrade it to master status using Partition Manager or DSREPAIR.

To restore replicas to the server that experienced the hard disk failure, you must complete the following steps. (These steps are explained in more detail on the following pages.)

1. Record the replicas and determine the master replica status.
2. Delete the Server and Volume objects.
3. Resolve any NDS errors and remove the replica pointers.
4. Install the hard drive and NetWare 4.
5. Place the replicas back on the server.
6. Restore the Directory data from a backup.
7. Confirm the correct bindery context.

Procedure: To restore the replicas to the server that experienced the hard disk crash, complete the following steps:

Record the replicas and determine the master replica status.

1. Document the replicas located on the server.

   To do this, use Partition Manager to select the Server object; then record the replicas listed.

2. If any replica displayed in Step 1 is a master replica, you will have to designate a new master on a different server in the tree.

   To do this, you will need to run DSREPAIR on a server that has an up-to-date read/write replica of the partition you need to change.

3. To change the replica type, follow these steps:

   a. At the server console, type

      LOAD DSREPAIR <Enter>

   b. Select **Advanced options menu**.

   c. Select **Replica and partition operations**.

   d. Select the partition you want to change.

   e. Select **Designate this server as the new master replica**.

   f. Press <Alt>+<F10> and <Enter> to exit DSREPAIR.

Delete the Server and Volume objects.

4. Using Partition Manager, delete the Server object of the server on which the hard drive crashed.

   Complete the following steps:

   a. Launch NetWare Administrator on a workstation.

   b. Select **Tools...Partition Manager**.

   c. Search the Directory tree to find the server you want to remove.

   d. Click on the server.

   e. Click on **Delete Server**.

      The Server object is also deleted on other servers that have replicas.

5. Using NETADMIN or NetWare Administrator, delete the Volume objects associated with the server.

Resolve any NDS errors and remove the replica pointers.

6. At the server console of the server that contains the copies of the partition in which the deleted Server object was located, type

   LOAD DSREPAIR <Enter>

7. Select **Replica synchronization** to show all partitions that existed on the deleted server.

8. Resolve all errors before continuing.

9. If a 625 error occurs while NDS is trying to connect to the server deleted in Partition Manager, wait to see if the error clears up. If the error remains, perform the steps listed below.

Normally, you need to perform the following steps only on the master replica. NDS should synchronize the deletion. If following the steps does not clear up the 625 errors, you will have to perform the same steps on all the servers containing replicas of the partition in question.

To remove the deleted server out of the replica pointer list, complete the following steps.

Use these procedures only to remove a deleted Server object. Using this procedure to try to delete a functioning server from a replica pointer list will corrupt the list.

a. At the server console, type

   LOAD DSREPAIR <Enter>

b. Select **Advanced options menu**.

c. Select **Replica and partition operations**.

d. Select the partition you want to edit.

e. Select **View replica ring**.

f. Select the deleted server name.

g. Select **Remove this server from the replica ring**.

h. Select **Yes**.

i. Press <Alt>+<F10> and <Enter> to exit DSREPAIR.

Install the hard drive and NetWare 4.

10. Install the new hard drive.

    Using procedures recommended by the manufacturer, verify that the hard drive is working properly. We recommend that the new hard drive have the same or larger storage capacity as the drive it replaces.

11. Reinstall NetWare the same way you previously installed it.

    After you have defined the configuration and copied the NetWare 4 files, the NDS portion of INSTALL appears on the screen. Install the server to the original context.

Place the replicas back on the server.

12. Use Partition Manager and the list you generated in Step 1 to place replicas on the server.

    The completion time needed to restore replicas depends on the speed of your LAN or WAN environment.

Restore the Directory data from a backup.

13. Use the backup system to restore the server's data.

    The backup system restores not only data but file ownership and user trustee information. If the backup system is unable to restore file ownership and trustees, you will have to recreate them manually.

Confirm the correct bindery context.

14. Set the bindery context (if needed).

    The restoration of the server and volume is complete.

Removing a Server from the Tree

On some occasions (if you are upgrading your hardware, for example), you may need to remove a server from the network.

Unlike other objects in the Directory tree, file servers cannot be deleted through the normal NETADMIN or NetWare Administrator interfaces. Servers require special consideration because they may contain replicas and other essential references to the Directory. These references must be deleted.

To remove a file server from the Directory tree, you must use either PARTMGR or Partition Manager from within NetWare Administrator.

Do not make partition changes while a server that will return is out of the Directory tree.

Removing a Server Using Partition Manager in NetWare Administrator

To remove a server using Partition Manager, complete the following steps in NetWare Administrator:

1. Select **Tools...Partition Manager**.
2. Search the Directory tree to find the server you want to remove.
3. Click on the server.
4. Click on **Delete Server**.

Removing a Server Using PARTMGR

1. At the prompt, type

   PARTMGR <Enter>

2. Search the Directory tree to find the server.
3. Press <**Del**>.
4. Confirm that you want to delete this object.

When a Server object is deleted, it is immediately removed from all replica rings.

Removing Volume Objects

You must remove any Volume objects that belonged to the server you removed. Use the normal delete procedures for Directory objects in NETADMIN or NetWare Administrator.

These volumes are represented with an Unknown Object icon after the server is removed.

Sending Updates to Synchronize Corrupted Replicas

Sending updates to all replicas should be your last resort if you are having trouble with your Directory. This means that you cannot log in or alleviate the problem using DSREPAIR. Sending updates is a nondestructive way to force all replicas to update to the master replica.

Use extreme caution with this function. This procedure updates all other replicas from the master replica of a partition, regardless of time stamp.

Using NetWare Administrator, complete the following steps to send updates to other replicas:

1. Using Partition Manager, select the partition in the Directory tree.
2. Click on **Replicas**.
3. Select the master replica.
4. Click on **Send Updates**.
5. Confirm the selection.

This procedure may cause excess traffic on your network because it sends the entire replica that is affected.

Summary

Detailed design of the Directory tree requires a coordinated plan of NDS security, partitions, replicas, and time synchronization. In the detailed design, you showed how your Directory tree would be affected by NDS security, where the partition boundaries would be located, which servers would have specific types of replicas from different partitions, and which time server types would be used for each server.

Taking some minor precautions now can save you time later. NetWare 4 provides a built-in protection mechanism for maintaining the security and accessibility of the Directory database. Proper partitioning and replicating of the database will ensure its integrity.

As the NetWare 4 network administrator, you are responsible for the following tasks:

- Minimizing the risk of Directory database inconsistencies
- Managing the downtime of services throughout your areas of responsibility
- Maintaining the integrity of the Directory database by determining if replicas are synchronized
- Taking decisive action to repair replicas that have become unusable

Notes

SECTION 6 Integrating and Managing NetWare 3

Introduction

In this section, you will integrate NetWare 3 with a NetWare 4 environment and provide manageability for NetWare 3 users, groups, and print services using NetWare Administrator within NetWare 4.

Objectives

At the end of this section, you will be able to do the following:

1. Integrate NetWare 3.1*x* with NetWare 4 using bindery services.
2. Manage NetWare 3.1*x* resources from NetWare Administrator using NetSync™.

Bindery Services

Bindery services is a feature of NetWare 4 that provides backward compatibility to pre-NetWare 4 clients and applications. Bindery services works by making objects in the hierarchical Directory appear as if they are in a flat *bindery.*

The Bindery

NetWare 2 and NetWare 3 servers have no directory services. Instead, they use a bindery as a database of objects. The bindery is not global in nature; therefore, each bindery-based server has its own separate bindery. The bindery contains objects known to and associated with the server, such as Users, Groups, Print Servers, and Print Queues.

The bindery does not contain many of the objects available in NDS™, such as Alias objects or Directory Map objects. Nor does it provide many of the object properties that are available with NDS. No hierarchy of objects exists in a bindery.

Since the bindery is server-based, bindery users are required to log in separately to each server that provides services they need.

Many network applications were written for NetWare 2 and NetWare 3 to provide network services. Print servers and data backup applications have been particularly popular. These applications make calls to the NetWare bindery to perform network functions.

Benefits of Bindery Services

Bindery Services allows companies that already have a large investment in NetWare 2 and NetWare 3 to use their bindery-based clients and applications until they can be upgraded to NDS. Once upgraded, clients and applications can then recognize and make use of NDS capabilities.

Bindery services associates NDS objects with a server to create a flat structure, or "bindery," for that server. This is done on a container-by-container basis. All objects in a container that is designated a server's *bindery context* are considered part of that server's "bindery."

All leaf objects within that container object can then be accessed both by NDS objects and by bindery clients, servers, and applications. Of course, only NDS objects that correspond to a bindery object are used by bindery services.

For more information on bindery services, see "Understanding Bindery Services" in the *Introduction to NetWare Directory Services.*

Bindery Context

The container object where you set a server's bindery services is called the *bindery context.* You can change the bindery context by using the SET command at the server console. You can set up to 16 bindery contexts for each server, so that all objects in those containers appear in the server's "bindery." The following figure illustrates bindery services when a bindery context is set for an Organizational Unit object.

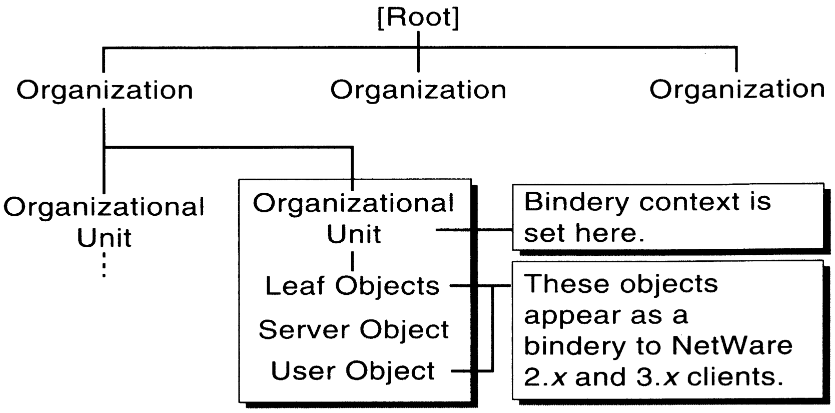

Figure 6-1: Bindery Context

Default Bindery Context for a New Server

When you install a NetWare 4 server in the Directory tree, a NetWare Server object is created in a container object. By default, the bindery context is set at that container object and bindery services is activated.

Although bindery services is enabled during installation of NetWare 4, you can change it (and the bindery context) with SET commands.

Naming Issues

When using bindery services with NetWare 4, you need to consider naming issues: bindery-compatible names and name conflicts.

Bindery-Compatible Names

The objects in the container object set as the bindery context must have bindery-compatible names. For example, the complete name for a User object might be

RussC.Corp.EMA

However, only the common name of the User object or RussC can be seen by bindery clients, applications, and servers using bindery services. Therefore, you should make sure that the object's common name conforms to the bindery naming rules.

This is one reason why you should avoid putting spaces in common names.

Name Conflicts

If you have the same common name in more than one container (see Figure 6-2), be aware of possible name conflicts. These conflicts occur when you set a server's bindery context to both containers. The first client to request the bindery service will receive the requested service; any other object with the same common name will be denied access.

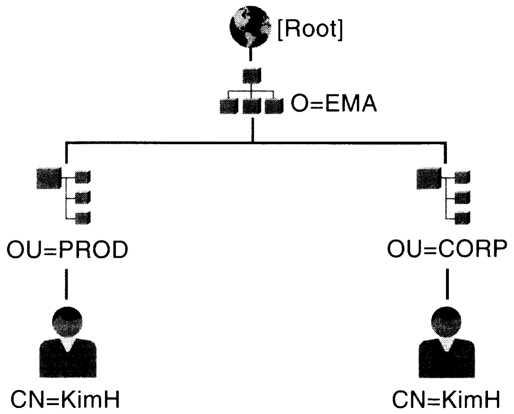

Figure 6-2: Naming Conflicts with Bindery Services

Bindery Services Considerations

By default, bindery services is enabled during the installation process. The bindery services context is set to the context in which the NetWare server is installed. If additional clients outside of this context need resources through bindery services, determine where access is needed; then set the bindery services contexts for those containers.

User objects for all users that log in to this NetWare 4 server from NetWare 3.*x* clients must be in the container where the server's bindery context is set.

Clients using the older shells (NETX) receive a BINDERY LOCKED error if they attempt to log in to a NetWare 4 server that does not have a bindery services context set.

Login scripts in NDS containers and User objects are not available through bindery services. If you need a bindery-based login to a NetWare 4 server, you must create and manage the bindery-based system login script and user login script with SYSCON (the NetWare 2 and NetWare 3 administrative utility).

Objects and properties, such as profile login scripts, that are unique to NDS are not available through bindery services.

When you create objects using bindery applications, all objects are placed in the Directory tree where you have set the server's bindery context.

A server must have a read/write or master replica of all partitions that include the server's bindery contexts.

Setting Bindery Services Context

You can set the bindery services context in one of the following ways:

- To set the context temporarily, type the following at the server console:

  SET BINDERY CONTEXT = ***context*** <Enter>

  Do not use a leading period (.) to specify context. Separate each context with a semicolon (;). For example, you would enter the following:

  SET BINDERY CONTEXT = OU=Prod.O=EMA;O=EMA

- To make a permanent change, change the SET BINDERY CONTEXT= line in the AUTOEXEC.NCF file to the new bindery services context; then restart the NetWare server.

- To make an immediate change and add the command to the AUTOEXEC.NCF file, use SERVMAN to change the SET BINDERY CONTEXT = command.

  1. Select **Server Parameters**.
  2. Select **Directory Services**.

     Bindery Context is the last command in the top window.

  3. Change the context to your new setting.
  4. Press <Esc> to exit.
  5. When you exit SERVMAN, you will be prompted to permanently save the change in the AUTOEXEC.NCF file.

Use the CONFIG command at the server console to determine the current bindery context.

For more information on bindery context, see "Understanding Bindery Services" in *Introduction to NetWare Directory Services,* SET in *Utilities Reference,* and "Bindery context path" in *Concepts.*

Hands-On Exercise 6-1: Using a NetWare 4 Server from a NETX Client

Activity: Access your NetWare 4 server as a NetWare 3.1*x* server using bindery services.

Procedure: Exit MS Windows and complete the following tasks.

1. Edit the workstation's CONFIG.SYS file to change the LASTDRIVE=Z statement to LASTDRIVE=E.
2. Edit your AUTOEXEC.BAT file to not load the workstation's connection and NetWare DOS Requester™ software when your workstation boots. Type **REM** before the STARTNET.BAT command in the AUTOEXEC.BAT file.
3. Reboot your workstation.
4. Load the DOS shell (NETX).

   a. Change to the drive and directory containing the DOS shell. Type

      C: <Enter>

      CD NWCLIENT <Enter>

   b. Run the client software. Type

      LSL <Enter>

      LAN driver <Enter>

      IPXODI <Enter>

      NETX <Enter>

5. Log in to the classroom NetWare 3.1*x* server. (The name of the server and username will be provided by your instructor.)

   a. Type

      F: <Enter>

   b. Type

      LOGIN *servername/username* <Enter>

   c. Enter the password, if required.

Logging in to the server sets your first search drive to SYS:\PUBLIC on the NetWare 3.1*x* server.

6. Using NetWare 3.1*x* utilities, attach to your NetWare 4 server as SUPERVISOR.

   a. Type

      ATTACH E*nnn*/SUPERVISOR <Enter>

   b. Enter the password you entered for Admin.

7. Using the NetWare 3.1*x* MAP utility, create network drive G: and map it to the SYS: volume of your NetWare 4 server. Type

   MAP G:=E*nnn*/SYS: <Enter>

8. Use NLIST to view NDS information from the PUBLIC directory on your NetWare 4 server.

   a. Type

      G: <Enter>

   b. Type

      CD \PUBLIC <Enter>

   c. From the PUBLIC directory of your NetWare 4 server, type the following:

      NLIST SERVER <Enter>

      NLIST VOLUME <Enter>

      NLIST USER <Enter>

      NLIST GROUP <Enter>

      NLIST PRINTER <Enter>

      NLIST QUEUE <Enter>

9. Which commands in Step 8c did not work and why?

10. Try to load NETADMIN. Type

    NETADMIN <Enter>

11. What happened and why?

12. Load the SYSCON utility from the NetWare 3.1*x* server. Type

    Z:SYSCON <Enter>

13. Select **User Information**.

14. View and manage the properties of the SUPERVISOR account.

    Which user account properties can you view and manage with SYSCON?

15. View and manage the properties of the Supervisor options.

    Which Supervisor options can you view and manage with SYSCON?

16. Enter the following commands in the system login script:

    WRITE "This is the Bindery Services System Login Script."

    MAP G:=*servername*\SYS:USERS\%LOGIN_NAME

17. Exit SYSCON.

18. Log in to the NetWare 3.1*x* classroom server as Supervisor and confirm your system login script. Type

    LOGIN *servername*/SUPERVISOR <Enter>

19. Edit your CONFIG.SYS and AUTOEXEC.BAT files to make them compatible with the NetWare DOS Requester.

    a. In CONFIG.SYS, change LASTDRIVE=E to LASTDRIVE=Z.

    b. In AUTOEXEC.BAT, remove REM from the command that calls STARTNET.BAT.

You have now completed the exercise.

Managing NetWare 3 from NetWare Administrator

NetSync is a NetWare loadable module that allows you to copy the binderies of multiple NetWare 3.1*x* servers into an NDS container on a NetWare 4 server. You can then manage the users, groups, and print services of the NetWare 3.1*x* servers as part of the Directory tree using the NetWare administrative utilities (NETADMIN and NetWare Administrator).

Using NetSync

You should use NetSync if

- You are running NetWare Name Service® (NNS) software and do not want to upgrade all servers in an NNS domain to NetWare 4 simultaneously.
- You need a temporary solution for central administration of a mixed network before full migration to NetWare 4.
- You want to make existing NetWare 3.1*x* users, groups, and print queues part of the NetWare 4 Directory without upgrading all NetWare 3.1*x* servers to NetWare 4.

Do not use NetSync if

- You have only NetWare 4 servers on your network.
- You intend to migrate all your NetWare 3.1*x* servers to NetWare 4 simultaneously or within a short time period.
- You plan to have a separate network supervisor account for each NetWare 3.1*x* server in your network from which someone will continue to manage users and groups with SYSCON and other NetWare 3.1*x* utilities.
- You do not need to make NetWare 3.1*x* users and groups part of the NetWare 4 Directory database.

For more information on NetSync, see *Installing and Using NetSync*.

How NetSync Works

To run NetSync, select one NetWare server to be the host for up to 12 NetWare 3.1*x* servers. You load NetSync NLM programs first on the NetWare host server and then on the NetWare 3.1*x* servers.

NetSync initially copies (uploads) all users and groups from the NetWare 3.1*x* servers' binderies into the bindery context of the NetWare server.

All bindery and NDS User and Group objects in that context are combined and copied to each NetWare 3.1*x* server. Each NetWare 3.1*x* server then has a copy of this "superbindery" (see Figure 6-3).

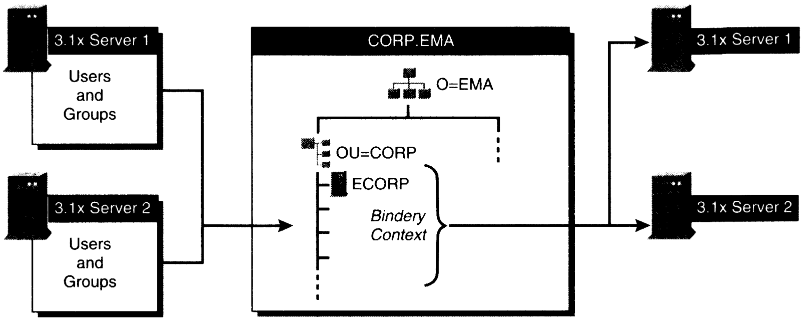

User and group objects from each 3.1x server are copied to the 4.1 server's.

All objects are copied back to each 3.1x server.

Figure 6-3: The NetSync Process

The configuration of a host NetWare server and multiple NetWare 3.1*x* servers is called the NetSync cluster. User and Group objects for all servers in the NetSync cluster can then be managed with the NetWare administrative utilities. Changes made to these objects are copied (downloaded) to the NetWare 3.1*x* servers in the cluster as they occur.

Once NetSync is installed, use only the NetWare administrative utilities, NETADMIN and NetWare Administrator, to manage user and group accounts. Changes made with SYSCON to the NetWare 3.1*x* servers in the cluster are not uploaded to the host server or to the other NetWare 3.1*x* servers in the cluster.

Installing NetSync

NetSync consists primarily of three NetWare loadable modules: NETSYNC4, NETSYNC3, and REMAPID.

NetWare Module	Description
NETSYNC4	Loaded on the NetWare 4.1 host server. Use this module to control the NetSync cluster.
NETSYNC3	Loaded on each NetWare 3.1*x* server in the NetSync cluster. This module copies (uploads) the NetWare 3.1*x* server's bindery information to the NetWare 4.1 server's bindery contexts; it then continues to communicate with the host server to receive updates to its bindery.
REMAPID	Autoloaded by NETSYNC3.NLM on every NetWare 3.1*x* server in the NetSync cluster. This module handles password synchronization and must remain loaded even if the NETSYNC3 module is unloaded.

Table 6-1: NetSync Modules

Each module should run continuously. Therefore, you should add the LOAD NETSYNC4 command to the NetWare 4 server's AUTOEXEC.NCF file, and the LOAD NETSYNC3 command to the NetWare 3.1*x* server's AUTOEXEC.NCF file.

NetSync has the following hardware requirements:

- At least one NetWare 4.1 server
- Up to twelve NetWare 3.1*x* servers per NetWare 4.1 server

NetSync has the following software requirements

- NetSync files
- One unused licensed connection on the NetWare 4.1 server

NetSync Installation Procedure

This procedure will be demonstrated by the instructor.

To prepare to install NetSync, do the following:

1. Resolve duplicate name conflicts.
2. Set the correct bindery context.

To install NetSync on a NetWare 4.1 server and copy files to NetWare 3.1*x* servers, do the following:

1. Load NETSYNC4 on a NetWare 4.1 server.
2. Enter the NetWare 3.1*x* server name.
3. Set the NETSYNC password.

> This password is used only to synchronize the servers. Do not use the same password as SUPERVISOR or ADMIN.

4. Copy the NetSync files to the NetWare 3.1*x* server.
5. Upload bindery data from a NetWare 3.1*x* server to a NetWare 4.1 server.
6. Enter the SUPERVISOR's username and password.

   The name of the NetWare 3.1*x* server appears in the authorized list.

   When prompted, add the new NETSYNC commands to the AUTOEXEC.NCF file.

To load NetSync on a NetWare 3.1*x* server, do the following:

7. Restart the NetWare 3.1*x* server to load the NetWare 4.1 files.

   The commands in the AUTOEXEC.NCF autoload NETSYNC3.NLM and REMAPID.NLM.

8. At the console prompt, enter the NetWare 4.1 server name you want this NetWare 3.1*x* server to synchronize with.
9. Enter the NetSync password you provided in Step 3.

Verify the NetWare 3.1*x* synchronization.

10. At the workstation, verify from NetWare Administrator that the NetWare 3.1*x* objects have been added to the first container in the bindery context.

    You might have to collapse and expand the container to see the changes.

11. Create a user from NetWare Administrator and verify that the user is created on each server in the server cluster.

12. Create Server and Volume objects for the 3.1*x* servers in the Directory tree.

Moving Print Servers

With NetSync, you can move all NetWare 3.1*x* print servers and merge them into a single print server on the NetWare 4 server. Existing print queues appear the same to users, but are serviced by the NetWare 4.1 print server. Your NetWare 3.1*x* printers are placed in the Directory as NetWare 4 Printer objects and made available to NetWare 4 clients.

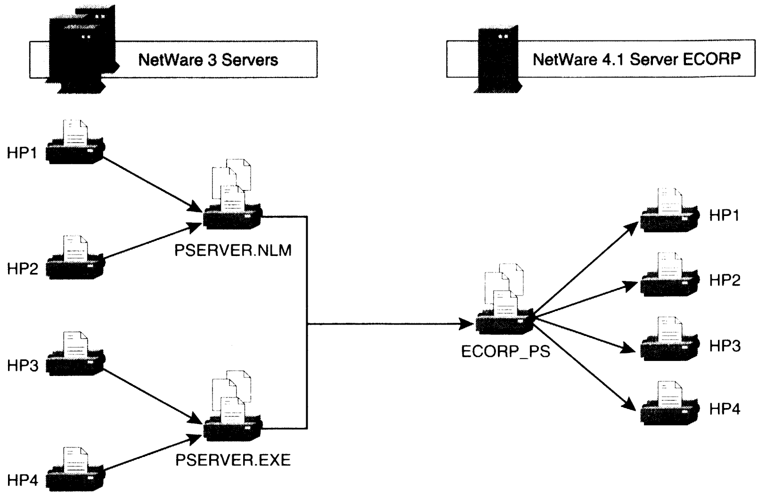

Figure 6-4: Moving Print Servers

The following changes related to network printing occur automatically to all NetWare 3.1*x* servers in the NetSync cluster:

- Print utilities are changed from NetWare 3.1*x* to NetWare 4.
- Databases for PRINTCON and PRINTDEF are upgraded to NetWare 4.

To merge your NetWare 3.1*x* print services with the NetWare 4 host server, complete the following steps:

1. Load NETSYNC4 on your NetWare 4 server.
2. Load NETSYNC3 on your NetWare 3 server.
3. Choose **Move a Print Server**.
4. Select the name of the print server you want to move.
5. In the NetWare 4 Directory, enter thє name of the print server you want to use.

If you type in an existing NetWare 4.1 print server, the NetWare 3 print server will be merged into the NetWare 4.1 print server. If you type in a new name, a new print server will be created in the Directory without a password.

6. Make new assignments for the Print Servers, Queues and Printers.

Summary

Bindery services allows your NetWare 4 server to appear as a NetWare 3.1*x* server for clients and applications that can only work with NetWare 3.1*x* servers. With bindery services, some NDS objects and properties can be seen, used, and managed as if they were part of a NetWare 3.1*x* bindery.

NetSync allows you to synchronize the binderies of multiple NetWare 3.1*x* servers with a single NetWare 4.1 server. Once synchronized, the user accounts, group accounts, and print services of the NetWare 3.1*x* servers in the NetSync cluster must be managed in the Directory tree as NDS objects using the NetWare administration utilities, NETADMIN or NetWare Administrator.

SECTION 7 Configuring NetWare 4 for Diverse Clients

Introduction

This section introduces the ability of NetWare 4 to support diverse clients. It also presents the procedure for installing NetWare Client™ for DOS and MS Windows on the client workstation and customizing the network connection with the NET.CFG file.

Objectives

Upon completion of this section, you will be able to do the following:

1. Configure the server to support diverse clients such as OS/2* and Macintosh*.
2. Install the NetWare Client for DOS and MS Windows using the NetWare Client installation software.
3. Create a custom network connection by modifying the NET.CFG file.

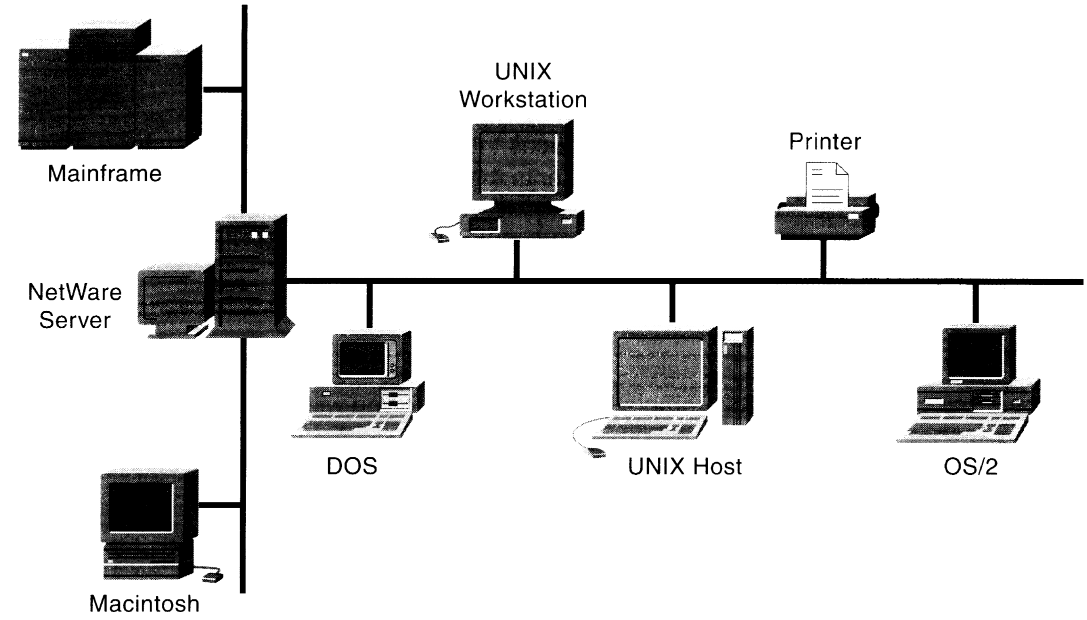

Figure 7-1

Configuring the Server for Diverse Clients

NetWare 4 can provide services to many kinds of clients, such as DOS, OS/2, Macintosh, and UNIX®. The main services NetWare 4 provides to diverse clients are communications, file services, and print services.

In this section, you will learn how to provide basic network services to Macintosh and OS/2 clients. Additionally, you will learn how to install the NetWare Client for DOS and MS Windows.

Providing network services to UNIX and other computers system is beyond the scope of this course. Communications support for UNIX clients is provided through the TCP/IP protocol. TCP/IP can be configured on a NetWare server with the INETCFG NLM, which is introduced in Section 8, "Managing Network Services." File system and print services for UNIX clients are provided by NetWare Networked File System (NFS), an optional product.

Configuring the Server for the Macintosh Client

NetWare 4 supports Macintosh workstations through NetWare for Macintosh and associated NLMs that support the AppleTalk Filing Protocol* (AFP). (AFP is the communication protocol used by Macintosh computers.) NetWare for Macintosh provides NetWare file services, print services, and routing services to Macintosh computers through its AFP support.

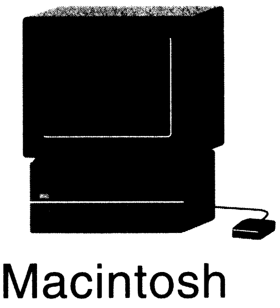

Figure 7-2: Macintosh Client

Three major functions are involved in setting up the server to service Macintosh clients:

- Installing NetWare for Macintosh on the server
- Configuring NetWare for Macintosh services
- Installing Macintosh client software

Installing NetWare for Macintosh

NetWare for Macintosh adds these capabilities to your NetWare network:

- Macintosh users can share files with non-Macintosh users, as well as other Macintosh users.
- Macintosh users can send print jobs to NetWare print queues.
- Non-Macintosh users can send print jobs to printers on an AppleTalk portion of the network.

NetWare for Macintosh can be installed during the initial NetWare 4 installation or after NetWare 4 has been installed.

The following installation steps can be completed *after* NetWare 4 has been installed.

1. Load INSTALL.NLM on the server. Type

   LOAD INSTALL <Enter>

2. Select **Product options**.
3. Select the **Choose an item or product listed above** option.
4. Select the **Install NetWare for Macintosh** option.
5. Specify the path to the NetWare for Macintosh software (you can accept the default path by pressing <Enter>).

   The window that opens specifies the path that INSTALL will use to locate the installation files. This is the path you specified during the most recent product installation.

6. If you are connected to a remote server, a window pops up for user authentication. Enter the username and the password.
7. To continue with the installation, select the **Install NW-MAC** option.

   A window displays messages indicating that the INSTALL utility is transferring the files. Copying the files may take a few minutes.

For more information, see "Installing NetWare for Macintosh" in *NetWare for Macintosh File and Print Services.*

The INSTALL utility creates a new directory in the SYS:SYSTEM directory called NW-MAC; below this directory, it creates the FONTS, PSUTILS, PPDS, ATPSCON, and SETUP subdirectories. The INSTALL utility copies the files from the source path you specified to the SYS:SYSTEM directory or the proper subdirectory.

8. Once the files are copied to the server, the Final Installation Options window opens. Each option enables you to configure a particular NetWare for Macintosh feature. The INSTALL utility does not implement any of your settings until you select Option 5 and press <Enter>.

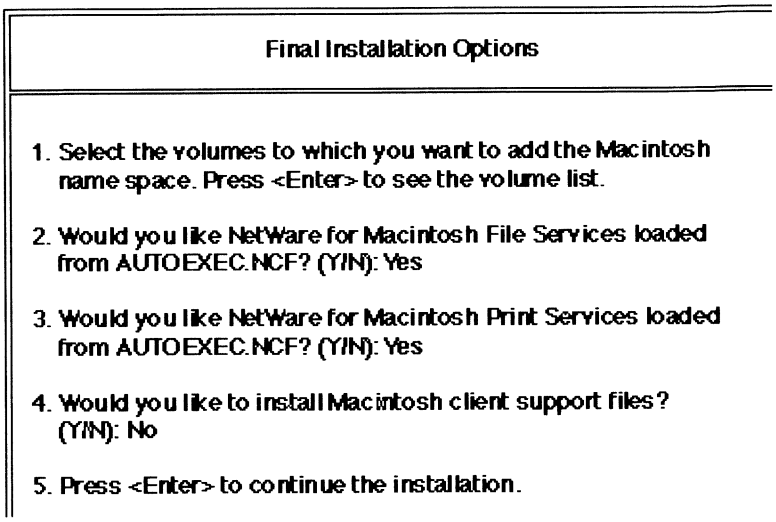

Final Installation Options

1. Select the volumes to which you want to add the Macintosh name space. Press <Enter> to see the volume list.

2. Would you like NetWare for Macintosh File Services loaded from AUTOEXEC.NCF? (Y/N): Yes

3. Would you like NetWare for Macintosh Print Services loaded from AUTOEXEC.NCF? (Y/N): Yes

4. Would you like to install Macintosh client support files? (Y/N): No

5. Press <Enter> to continue the installation.

Figure 7-3: Final Installation Options

a. Option 1 refers to adding name space to volumes on the server.

   Press <Enter> to see the volumes on the server.

b. Select a volume by pressing <Enter>; this marks the volume with an "X".

You must load the Macintosh name space on the SYS: volume in order to install the Macintosh client software.

c. Escape out of this window and press <Enter> to the **Yes, save changes and continue** prompt.

For more detailed installation steps, review "Installing NetWare for Macintosh" in *NetWare for Macintosh File and Print Services.*

d. Option 2 refers to adding the appropriate commands to the AUTOEXEC.NCF file so that the Macintosh software will load automatically when the server starts up. Select **Yes** and press <Enter>.

e. Option 3 refers to adding the appropriate commands to the AUTOEXEC.NCF file so that the Macintosh print services software loads automatically when the server starts up. Select **Yes** and press <Enter>.

f. Option 4 refers to installing the Macintosh client support files. The client support files allow Macintosh users to install the necessary software on to their workstations from the server. Select **Yes** and press <Enter>.

g. Option 5 accesses the confirmation window when you press <Enter>. Select **Yes** and press <Enter>.

After the NetWare for Macintosh installation is complete, the NetWare for Macintosh Configuration window appears.

Configuring NetWare for Macintosh Services

Once the installation is complete, the NetWare for Macintosh services must be configured and enabled. The following tasks should be completed to enable file and print services:

- Configure and enable AppleTalk.
- Customize file services.
- Configure and enable print services.

Each of these tasks can be accomplished from the NetWare for Macintosh Configuration window that appears once the NetWare for Macintosh installation is complete.

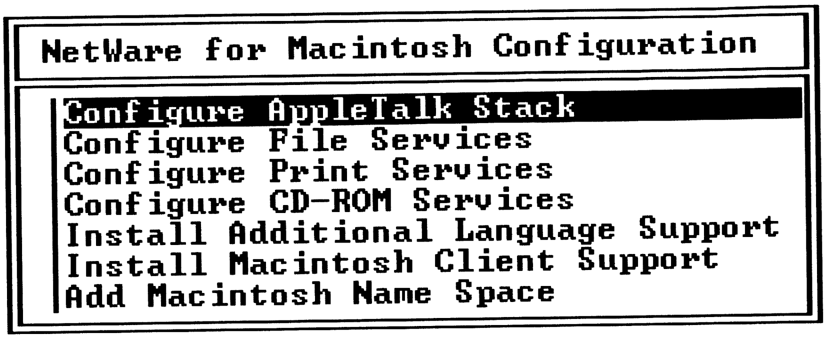

NetWare for Macintosh Configuration

Configure AppleTalk Stack
Configure File Services
Configure Print Services
Configure CD-ROM Services
Install Additional Language Support
Install Macintosh Client Support
Add Macintosh Name Space

Figure 7-4: NetWare for Macintosh Configuration Window

When you make a selection from this window, the INSTALL utility loads individual NLMs or allows you to install additional files and name space. The following table identifies the NLMs that load from the NetWare for Macintosh Configuration window.

Command Item	NLM
Configure AppleTalk Stack	INETCFG.NLM
Configure File Services	AFPCON.NLM
Configure Print Services	ATPSCON.NLM
Configure CD-ROM Services	HFSCDCON.NLM

Table 7-1: Configuration NLMs

Each NLM can be loaded at the server prompt as well as from the NetWare for Macintosh Configuration window.

Configuring and Enabling AppleTalk

Before the Macintosh clients can communicate with the NetWare server, you must use INETCFG to do the following:

- Identify the network board.
- Enable AppleTalk protocol routing.
- Bind the AppleTalk protocol to the LAN driver.

Customizing File Services

NetWare file services are immediately available to Macintosh clients as soon as NetWare for Macintosh is installed on the server.

With AFPCON, you can customize NetWare for Macintosh file services to enhance performance and maximize your network's file services capabilities. The following is a partial list of custom file services provided by AFPCON:

- Allow Guest Logins
- Allow User to Save Password for Auto Logins
- Set Maximum Number of AFP Connections
- Shut Down AFP Server
- Restart AFP Server
- Volume Information

For more information on configuring AppleTalk, see "Example 1: Basic Configuration" under "Network Configuration Examples" in *NetWare 4.1 AppleTalk Reference.*

For more information on NetWare for Macintosh file services, see "Managing AppleTalk Print Services" in *NetWare for Macintosh File and Print Services.*

Configuring and Enabling Print Services

When you install NetWare for Macintosh, the INSTALL program automatically copies to the server all files associated with NetWare for Macintosh print services. After installation, you must use ATPSCON to configure and enable NetWare for Macintosh print services.

The following print services setup and configuration functions can be performed with ATPSCON:

- Quick Configuration
- Configure Printer Servers
- Configure Spoolers
- Define Printer Models
- Log Options
- Management Options
- Change Context

NetWare for Macintosh print services provides bidirectional communication with printers. This bidirectional communication makes AppleTalk a preferred protocol for managing advanced printing environments.

For more information on NetWare for Macintosh print services, see "Planning AppleTalk Print Services" and "Setting Up AppleTalk Print Services" in *NetWare for Macintosh File and Print Services.*

Accessing the NetWare for Macintosh Configuration Window

After the initial installation of NetWare for Macintosh, you may need to access the NetWare for Macintosh Configuration window again. Do the following to access the configuration window:

1. At the server prompt, type

   LOAD INSTALL <Enter>

2. Select **Product options**.
3. Select **View/Configure/Remove installed products**.
4. Select **NW-MAC** from the Currently Installed Products window.

Installing Macintosh Client Software

Once NetWare for Macintosh is installed on the NetWare server, you can log in to the server from your Macintosh workstation (using AppleShare) and receive limited access to network services. To take full advantage of the NetWare 4 network, you should install the NetWare 4 client software (MacNDS) on each Macintosh workstation.

The MacNDS client software for the NetWare 4 operating system provides access to NetWare Directory Services for Macintosh workstations running under the System 7 operating system.

During the NetWare for Macintosh installation, you can copy the MacNDS client files onto the NetWare server (see Step 8f under "Installing NetWare for Macintosh" in this section). The Macintosh client files are stored on the server in the form of a self-extracting archive (MacNDS.SEA) file.

To install the Macintosh client, do the following:

1. Log in to the server containing the installation files.
   a. Enable bindery services on the server containing the installation files.
   a. Connect to the server.
   b. Log in as user Supervisor.
   c. Access the volume on which you installed the client installation files.
2. Locate the correct self-extracting archive in the PUBLIC\MAC folder.
   a. Locate the proper language folder.
   b. Locate the MacNDS.SEA file in the language folder.
3. Double-click on the MacNDS.SEA file.
4. In the dialog box that appears, identify the location where the files should be installed.

   The default setting will create a MacNDS folder within your language folder.

For information about installing the Macintosh client software, see "Installing the MacNDS Client Software" in *Using MacNDS Client for NetWare 4.*

Configuring the Server for the OS/2 Client

Once communication and print services have been set up for DOS workstations, no additional server configuration is needed to provide file and print services to an OS/2 client.

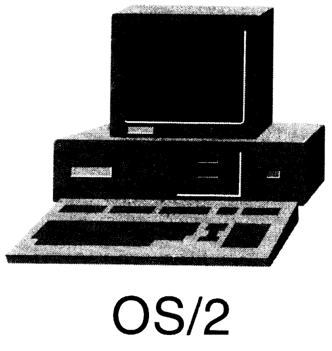

Figure 7-5: OS/2 Client

The NetWare 4 server communicates with an OS/2 client using the same IPX/SPX protocol used to communicate with DOS workstations. NetWare file and print services configured for DOS workstations can be used by OS/2 clients.

While OS/2 clients can use the DOS file system, OS/2 uses long names and extended attributes. *Long names* are filenames that can be longer than those used by DOS. (DOS restricts a filename to eight characters plus a three-character extension.) Long names can be up to 255 bytes. OS/2 *extended attributes* are associated attributes (names and values describing the file) attached to the long names.

For this reason, the OS/2 name space should be added to a NetWare volume to provide OS/2 extended file services to OS/2 clients.

Two major functions are involved in providing basic network services to OS/2 clients:

- Configuring the file system for OS/2 clients
- Installing OS/2 client software

Configuring the File System for OS/2 Clients

The NetWare 4 name space for OS/2 installs the High-Performance File System (HPFS). HPFS replaces the FAT file system used by DOS. HPFS provides compatibility with OS/2 and DOS files.

To initially configure a volume for a new name space, you must first load the name space module (.NAM) and then use the ADD NAME SPACE command to configure the volume. For example, to configure the SYS: volume to contain OS/2 name space, you would execute the following commands at the server console:

```
LOAD OS2 <Enter>
ADD NAME SPACE OS2 TO SYS <Enter>
```

The LOAD OS2 command identifies the name space and loads the protocols necessary to use the name space. The ADD NAME SPACE command creates the name space on the volume. It creates entries in the Directory table based on the directory and file naming conventions of the OS/2 file system.

Execute ADD NAME SPACE only once during initial configuration of the volume. After a volume is configured for a new name space, you only need to load the name space module NLM.

After the name space has been added to a volume, the corresponding name space module autoloads each time you mount a volume.

Installing OS/2 Client Software

The NetWare Client for OS/2 software enables OS/2 workstations to access NetWare servers. After you install the NetWare Client for OS/2, you can connect to a NetWare network and perform basic network tasks, such as accessing NetWare file systems and redirecting print jobs to NetWare print services.

The NetWare Client for OS/2 installation software can be found on the installation diskettes, the NetWare 4 installation CD-ROM, or in the SYS:PUBLIC\CLIENT\OS2 directory of a NetWare server where these files have been installed.

Installation diskettes can be create by executing the MAKEDISK software found on the NetWare 4 installation CD-ROM or in the SYS:PUBLIC\CLIENT\OS2 directory.

Details on installing and using NetWare Client for OS/2 are found in *NetWare Client for OS/2 User Guide.*

Installing DOS Connection and Client Software

You can easily install the Client software for DOS and MS Windows and set up the appropriate configuration files using the NetWare Client installation software.

Workstation Hardware and Software Requirements

The workstation hardware and software requirements are as follows:

- IBM PC or compatible
- XT, AT, 8088, 286, 386, 486 (SX, SLC, DX, etc.), or higher processor
- DOS and MS Windows (optional)
- 4 MB of disk space (5 MB if MS Windows software is installed)

The NetWare Client installation program sends an error message warning you that you have insufficient disk space if the total available disk space is below 5 MB when you are installing files for both DOS and MS Windows.

Network Board Configuration

The following table identifies how you can access the network board information required by the NetWare Client installation process.

If you have	Then
EISA or MCA network boards	Run the workstation's setup or reference program. This program lists the values for your network board settings.
ISA network boards	Look at the network board to obtain the specific settings. The documentation provided with your network board should tell you where to find each setting value.
PCI Local Bus network boards	Run the workstation's setup or reference program. This program lists the values for your network board settings.
Note: If you already have a network connection, change to the PUBLIC directory. Then type **NVER** and press <Enter>.	

Table 7-2: Network Board Configuration

NetWare Client Installation Procedure

The NetWare Client installation software is menu driven and prompts the user for all pertinent information. Figure 7-6: shows the NetWare Client Install screen.

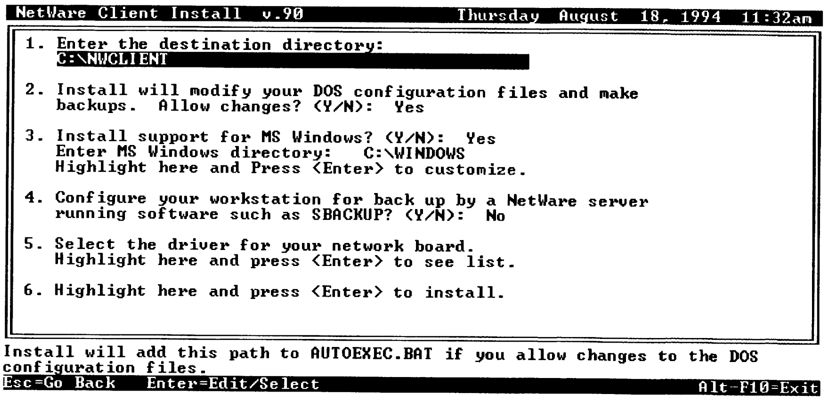

```
 NetWare Client Install  v.90                     Thursday  August  18, 1994  11:32am

  1. Enter the destination directory:
     C:\NWCLIENT

  2. Install will modify your DOS configuration files and make
     backups.  Allow changes? (Y/N):  Yes

  3. Install support for MS Windows? (Y/N):  Yes
     Enter MS Windows directory:   C:\WINDOWS
     Highlight here and Press <Enter> to customize.

  4. Configure your workstation for back up by a NetWare server
     running software such as SBACKUP? (Y/N):  No

  5. Select the driver for your network board.
     Highlight here and press <Enter> to see list.

  6. Highlight here and press <Enter> to install.

Install will add this path to AUTOEXEC.BAT if you allow changes to the DOS
configuration files.
Esc=Go Back    Enter=Edit/Select                                      Alt-F10=Exit
```

Figure 7-6: NetWare Client Install Screen

The following table identifies the installation steps for installing a client from either diskettes or a network directory.

Installing from diskettes	Installing from a network directory
1. Insert the *NetWare Client for DOS and MS Windows Disk 1* into a disk drive.	1. From the SYS:PUBLIC directory, change to the CLIENT/DOSWIN subdirectory. For example, type **CD CLIENT/DOSWIN** <Enter>
2. Change your drive to the drive that contains the diskette. For example, type **A:** <Enter>	2. Type **INSTALL** <Enter>
3. Type **INSTALL** <Enter>	3. Follow the instructions on your screen.
4. Follow the instructions on your screen.	

Table 7-3: Installation Procedure

For details on installing the DOS and MS Windows clients, see “Installing or Upgrading NetWare Client Software” in *NetWare Client for DOS and MS Windows User Guide.*

Customizing the Network Connection

As part of the installation process, certain information is placed in the NET.CFG file. This information automates the user's connection and login process.

You may need to customize the user's NET.CFG files for the user's network environment. One reason to modify the NET.CFG file is to automatically attach the user to a specific server and to place the user in the proper context in the Directory tree.

The following commands should be placed in the NetWare DOS Requester™ section of each user's NET.CFG file:

```
NAME CONTEXT= "Directory_tree_context"

PREFERRED SERVER= server_name

or

PREFERRED TREE= tree_name
```

Here is an example of a customized NET.CFG file:

```
NETWARE DOS REQUESTER
    FIRST NETWORK DRIVE = F
    NAME CONTEXT= "OU=CORP.O=EMA"
    PREFERRED TREE= EMA_TREE
```

Customized NET.CFG files can be created in two ways:

- Editing the NET.CFG file after installing the client
- Editing the INSTALL.CFG file before installing the client

Editing the NET.CFG File after Installing the Client

After installing the NetWare Client for DOS and MS Windows, you will need to customize each user's NET.CFG file with additional options. You can use any DOS text editor to edit the NET.CFG file.

Editing the INSTALL.CFG File before Installing the Client

You may need to set up several users whose NET.CFG files use the same NET.CFG settings (such as the PREFERRED SERVER, PREFERRED TREE, or NAME CONTEXT options described above). Instead of editing each workstation's NET.CFG file, you can change the INSTALL.CFG file. As part of the installation, the options in the INSTALL.CFG file are used to create the NET.CFG file on each workstation.

When the client software has been installed on a NetWare 4 server, the INSTALL.CFG file can be found in the \PUBLIC\CLIENT\DOSWIN directory. You can use any DOS text editor to change the NetWare DOS Requester settings in the INSTALL.CFG file.

Summary

The NetWare for Macintosh software enables a NetWare server to provide network communication, file, and print services to Macintosh clients. The Macintosh client software should be installed on each Macintosh workstation to take full advantage of the services provided by NetWare 4.

OS/2 clients using the NetWare Client for OS/2 can immediately use the same communication, file, and print services provided for DOS clients by a NetWare server. You should install the OS/2 name space on any NetWare volumes that will store OS/2 files. This allows the OS/2 client to use long filenames and OS/2 extended attributes.

The DOS and MS Windows clients can be easily installed with the NetWare Client installation software. It will modify all necessary configuration files. You can create custom NET.CFG files during client installation by entering NET.CFG parameters in the INSTALL.CFG file before performing the client installation.

Notes

SECTION 8 Managing Network Services

Introduction

This section introduces you to supplementary services provided with NetWare 4 concepts and procedures about NetWare MultiProtocol Router™, message handling, and language support in a NetWare 4 environment.

Objectives

At the end of this section, you will be able to do the following:

1. Define the purpose of the NetWare MultiProtocol Router (MPR) 3.0 and give an example of its function within a network.
2. Determine appropriate use of messaging services through MHS Services for NetWare 4®.
3. Create a Distribution List object and assign members.
4. Assign mailboxes to User, Group, Distribution List, and Organizational Role objects.
5. Send and receive E-mail with FirstMail™.
6. Enable internationalization.

NetWare 4 Internetworking Software

NetWare 4 includes a suite of software that provides internetworking between dissimilar types of systems. This suite runs on standard PC-compatible hardware, and allows you to integrate a variety of LANs, regardless of the transport protocols being used. This software is called the NetWare MultiProtocol Router (MPR).

NetWare MPR provides the following services:

- It connects dissimilar media, frame types, and transports at all levels in the ODI layers.
- It off-loads routing tasks from busy servers.
- It provides security by isolating traffic on LAN segments.

You should manage the configuration of transport protocols on the network using a NetWare utility called INETCFG.NLM. Use INETCFG.NLM to configure and manage MPR. The next few pages demonstrate how to set up and manage multiple protocols on your network.

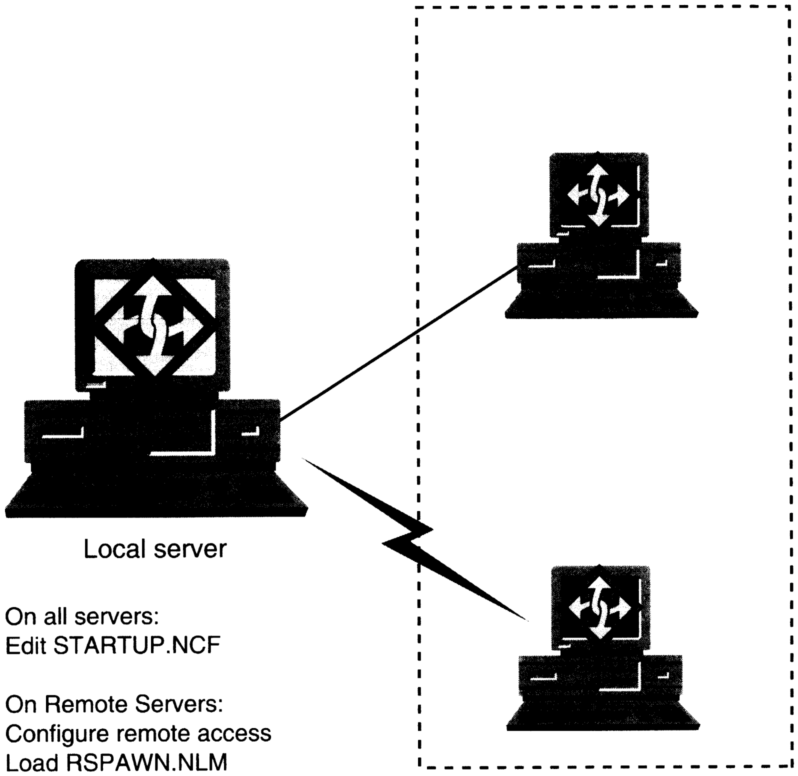

Figure 8-1

NetWare MPR 3.0

NetWare MPR 3.0, like other Novell products, is installed at the server console with INSTALL.NLM. The installation procedures are described in the *NetWare MultiProtocol Router 3.0 Installation and Basic Configuration Guide.*

For more detailed information on MPR and interoperability issues, refer to Novell Course 740, *Internetworking with NetWare MultiProtocol Router.*

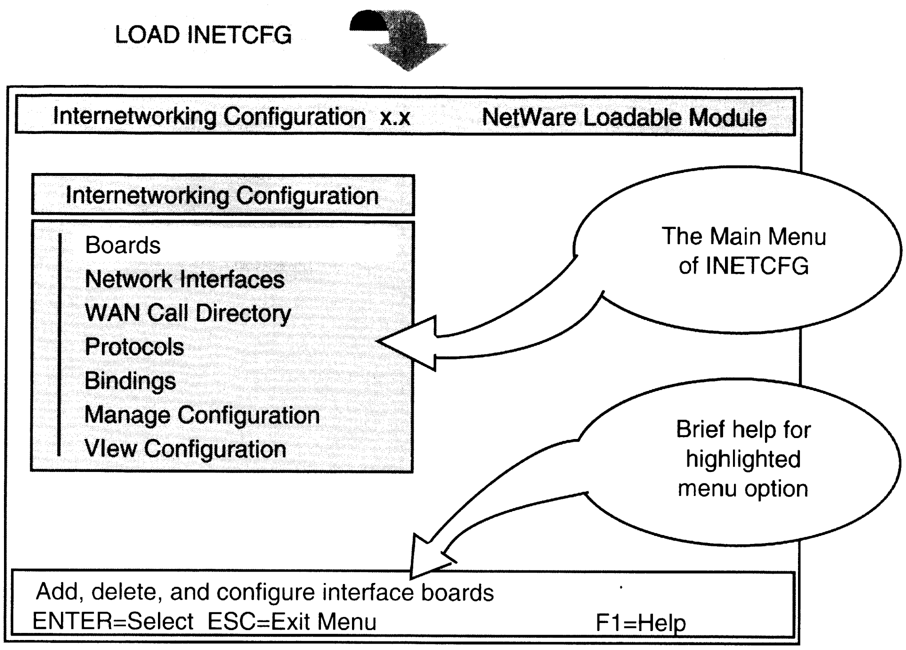

Add, delete, and configure interface boards

Figure 8-2

Using INETCFG

After NetWare MPR 3.0 is installed, you can begin configuring routing and bridging capabilities. To do this, you will use the INETCFG.NLM. INTECFG is a menu-driven utility that simplifies routing and bridging configuration while providing an extensive list of configurable parameters.

Loading the INETCFG Utility

When you load INTECFG for the first time, it does the following:

1. Creates a new AUTOEXEC.NCF file.

   INETCFG saves your old AUTOEXEC.NCF file to the backup file SYS:\SYSTEM/AUTOEXEC.BAK; if you must restore this file later, copy it to AUTOEXEC.NCF and restart the server.

2. Prompts you to import the LOAD and BIND commands from the old AUTOEXEC.NCF to the new file.

   After you have transferred commands to the new AUTOEXEC.NCF file, you should perform any future edits through INETCFG.

3. Prompts you to transfer LAN driver, protocol, and remote access commands from the old AUTOEXEC.NCF file to the new file. Answering Yes to this prompt allows INTECFG to automatically transfer these commands for you.

Some application services (such as NFS, AFP, and ATPS) still need to be configured in AUTOEXEC.NCF using LOAD commands. Refer to the *NetWare MultiProtocol Router 3.0 Advanced Configuration and Management Guide* for more information.

4. Inserts three commands in the new AUTOEXEC.NCF file in the following order:

   - LOAD CONLOG – Enables the server to begin a log file of the messages that appear during system initialization. The messages are stored in the ETC/CONSOLE.LOG file.

     The CONLOG.NLM is discussed in additional detail later in this section.

   - INITSYS.NCF – Enables the server to initialize using the configuration information in the INETCFG database.

   - UNLOAD CONLOG –Stops the logging of console messages when startup is complete.

Navigating the INETCFG Main Menu

The INETCFG main menu appears in Figure 8-3.

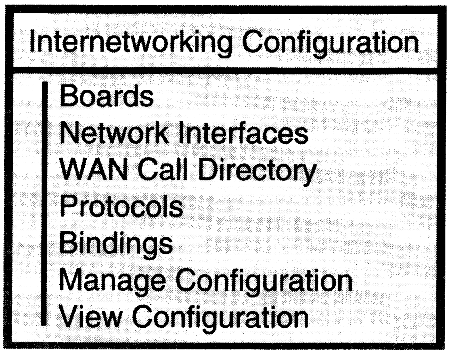

Internetworking Configuration

Boards
Network Interfaces
WAN Call Directory
Protocols
Bindings
Manage Configuration
View Configuration

Figure 8-3: INETCFG Main Menu

In general, the order in which these menu options are selected during configuration is top to bottom. The menu options are organized in a logical way; first you configure your boards; then you configure your protocols. Then you bind your protocols to the boards.

A brief description of each menu option appears below.

Boards Menu Option

The Boards menu option is used to configure LAN and WAN boards. Configuring the boards in INETCFG builds the correct LOAD commands in the AUTOEXEC.NCF file.

Network Interfaces Menu Option

The Network Interfaces option is used to configure ports on multiport boards. You can configure each port on a multiport WAN board to use a different WAN protocol, such as PPP, Frame Relay, or X.25.

WAN Call Directory Menu Option

The WAN CALL Directory option is used to define WAN call destinations for PPP (the WAN options in the WAN • Extensions and SNA • Extensions product).

Protocols Menu Option

The Protocols option is used to configure routing protocol operation, such as IPX, RIP, and NLSP. Configuring the protocols in INETCFG builds the correct LOAD commands for the protocols in the AUTOEXEC.NCF file.

Bindings Menu Option

The Bindings option is used to bind the enabled protocols to configured interfaces. The Binding option enables you to bind protocols on a per-interface basis. For example, on one interface you can bind IPX RIP and on another interface you can bind NLSP. Configuring the bindings in INETCFG builds the correct BIND commands in the AUTOEXEC.NCF file.

Manage Configuration Menu Option

The Manage Configuration option is used to configure SNMP parameters and remote access to the server. In addition, you can edit the AUTOEXEC.NCF file from this menu option.

View Configuration Menu Option

The View Configuration option is used to view LOAD and BIND commands and view console messages that were logged during system initialization.

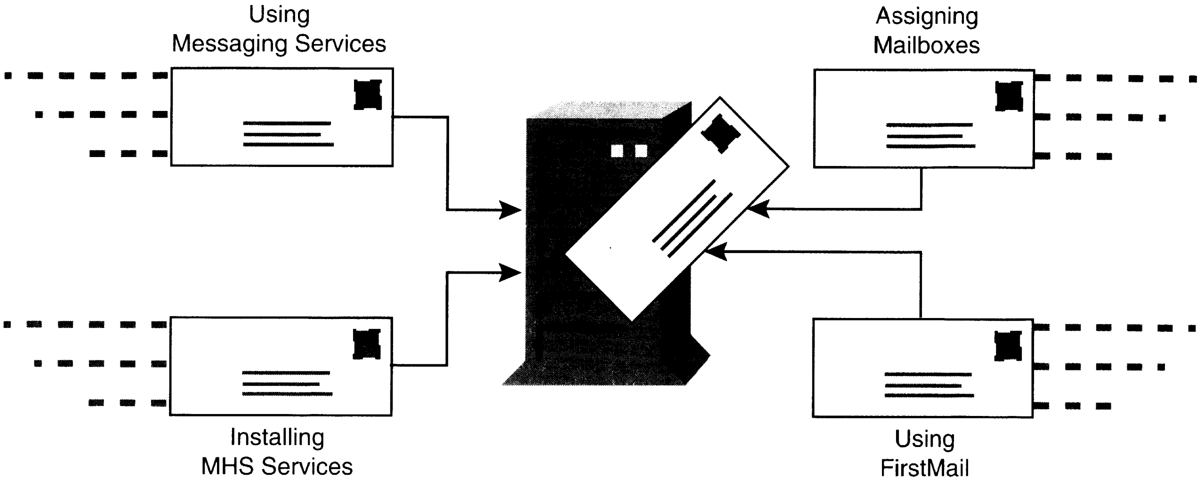

Figure 8-4

Messaging Services

This section presents the following topics:

- Using messaging services
- Installing MHS Services for NetWare 4
- Assigning mailboxes to NDS objects
- Using FirstMail to send and receive E-mail

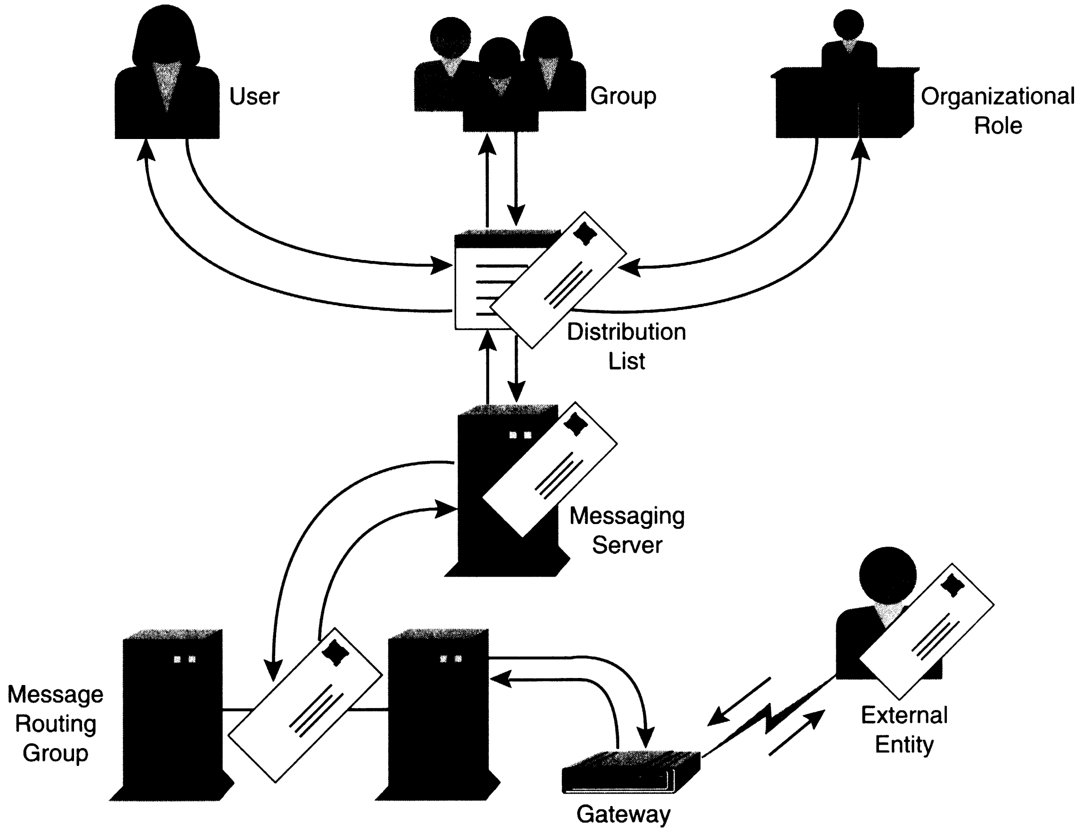

Figure 8-5

Using Messaging Services

Messaging services include storing, accessing, and delivering text, binary, graphic, digitized video, and audio data.

Messaging services are similar to file services. However, unlike file services, messaging services deal actively with the communication interactions between computer users. Instead of simply storing data files, messaging services transport data from point to point and notify the user of awaiting messages.

Installing MHS Services for NetWare 4

To install MHS Services for NetWare 4, you should perform the following tasks:

1. Identify your hardware and software requirements.
2. Install MHS Services during or after your initial NetWare installation.

Hardware Requirements

The MHS Services software is provided on the NetWare 4 installation CD-ROM.

MHS Services for NetWare 4 has no mailbox limit, meaning that you may identify as many user mailboxes as you wish. However, the number of concurrent users is limited by your NetWare 4 license. If additional concurrent users are required, your NetWare license must be upgraded.

To install MHS Services on a NetWare 4 server, the server must have the following hardware resources available:

- 500 KB of available RAM
- 2.5 MB of disk space for program storage, plus additional disk space for user mailboxes

The amount of disk space required for user mailboxes will vary, depending on the size and number of messages that users store in their mailboxes.

The NetWare 4 server performing as a messaging server should meet the following minimum criteria:

- 386 class (or better) processor
- 12 MB of RAM
- 65MB hard disk
- CD-ROM or mapped network connection to installation files directory

The preceding minimum requirements are not sufficient to handle more than 10 users or more than 100 messages per day. Therefore, the following additional requirements are suggested for messaging engines that process more than 100 messages per day:

- 16 MB of RAM plus any additional RAM required to maintain more than 30% free cache buffers (see the note that follows)
- 65 MB of hard disk space (for the default NetWare 4 installation) plus an additional 5 MB per mailbox (User, Distribution List, and so on)

The standard Novell algorithm for calculating required RAM is found in the *Installation and Upgrade* manual.

Installation Procedures

The following options are available for installing MHS Services:

- MHS installation after the NetWare 4 server is installed
- MHS installation as a part of NetWare 4 server installation

Installing MHS Services after the NetWare 4 Server Is Installed

To add MHS Services after the NetWare 4 server has been installed, complete the following procedures:

1. At the server console, type

   LOAD INSTALL <Enter>

2. Select **Product options**.
3. Select **Choose an item or product listed above**.
4. Select **Install NetWare MHS**.
5. Specify the path to the MHS Services software (you can accept the default path by pressing <Enter>).
6. In the Name field of the Postmaster General Authentication window, enter the full distinguished name (example: Admin.CORP) of the user who will be Postmaster General (the Postmaster General is typically Admin).
7. In the Password field, enter the Postmaster General's password.
8. Press <Enter> to continue.
9. If the NetWare server on which you are installing MHS Services has multiple volumes, select the volume on which you want to install the MHS Services database.
10. Press <Enter> to continue installing optional products, or exit from the INSTALL program.
11. Add the **LOAD MHS** command to the server's AUTOEXEC.NCF file.

Installing MHS Services during the Initial NetWare 4.1 Installation

Installing MHS Services during the initial NetWare 4.1 installation is very similar to adding MHS after NetWare 4 is installed. Replace Step 1 in "Installing MHS Services after the NetWare 4 Server Is Installed" (on the previous page) with the following step:

1. Install NetWare using the **Customized installation of NetWare 4.1** option.

   When the operating system installation is complete, the Other Installation Actions menu appears.

Continue with Steps 2 through 11 from the previous page.

For further information, refer to "Installing and Running MHS Services" in *MHS Services for NetWare 4.*

What the MHS Installation Creates

When you install MHS Services during an initial NetWare 4 installation, the following events occur:

- An NDS object, the Message Routing Group, is created (the default name is MHS_ROUTING_GROUP).
- The default owner (Admin) of the Message Routing Group is assigned.
- The default Messaging Server is created.
- A Postmaster for the Messaging Server is assigned.
- The default Mailbox Location for Admin is assigned to the Messaging Server and to the default Mailbox ID.
- The default Messaging Server is added to the default Message Routing Group; the default Message Routing Group is assigned to the Message Routing Group property of the Messaging Server.
- The local NetWare server's Messaging Server property is assigned the default Messaging Server; the local NetWare server is assigned to the NetWare Server property of the default Messaging Server.
- FirstMail for DOS and FirstMail for MS Windows are automatically installed in the SYS:\PUBLIC directory.

Among the many files and subdirectories that are created, \MHS\MAIL\USERS\username\FIRST.APP is the default path to a user's mailbox. This mailbox location is linked through NDS to the corresponding Messaging Server object, which is created during MHS Services installation.

If you delete the Admin User object, the owner of the Message Routing Group and the Postmaster of the Messaging Server will need to be reassigned. Failure to do so will cause extensive troubleshooting problems.

Assigning Mailboxes to NDS Objects

You can use E-mail and messaging applications to send MHS messages to any NDS object that can be assigned a mailbox. In addition to understanding how to assign mailboxes, you should familiarize yourself with the general purpose of the NDS objects that are used by MHS Services.

The following NDS objects exist outside of MHS Services but require MHS-specific property pages:

- User
- Group
- Organizational Role
- Organizational Unit

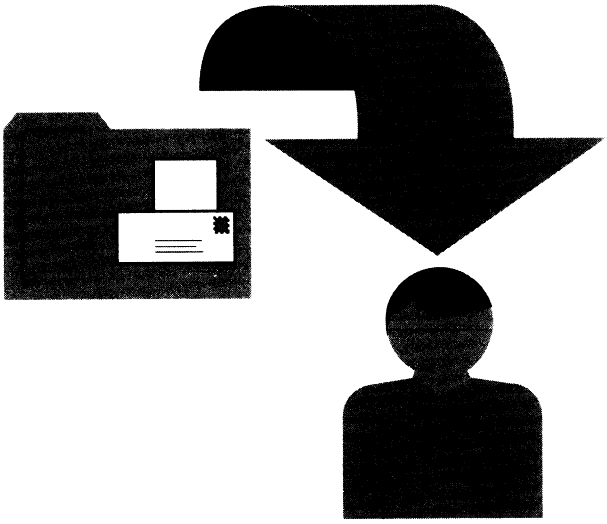

Figure 8-6: Assigning Mailboxes to NDS Objects

For further information, see "Managing MHS Mailboxes" in *MHS Services for NetWare 4*.

The following NDS objects are specifically designed for use with MHS Services:

- Messaging Server
- Message Routing Group
- Distribution List
- External Entity

Messaging Server

The Messaging Server object enables messaging services. This object identifies the location of the message directory structure (\MHS) serviced by the messaging engine. It also identifies the NetWare server responsible for running the messaging engine.

During a basic installation, this object is created automatically and does not require you to configure anything.

Message Routing Group

The Message Routing Group object represents a group or cluster of messaging engines that communicate directly with each other for transferring messages. A default Message Routing Group is created during the first installation of MHS Services. Subsequent installations of other MHS Services messaging engines will default to this group.

If you want multiple MHS Services engines to route messages to each other, make sure that all the Message Routing Groups are assigned the same name in the Name property.

Distribution List

The Distribution List object can be used to address mail to multiple mailboxes. It also reduces the network traffic for messages to mailboxes on multiple Messaging Servers. Only one message copy is delivered to a Distribution List mailbox; the message is then replicated for every mailbox on the Distribution List.

Group objects can also be used to address mail to multiple user mailboxes; however, a group's membership cannot contain other groups. Unlike the Group object, a Distribution List object is an NDS object created to simplify the addressing of messages. Distribution List object members do not share a login script or trustees assignments. A Distribution List's membership can be nested; in other words, it can contain other Distribution Lists.

Use the following procedure to create a Distribution List:

1. Using NetWare Administrator, select **Create** from the Object menu.
2. Select the **Distribution List** object and click on **OK**.
3. Enter a unique Distribution List Name.
4. Click on the **Browse** button to the right of the Mailbox Location field.
5. Select the appropriate Messaging Server object from the Objects list box (choose the Messaging Server on which the majority of the Distribution List members have mailboxes) and click on **OK**.
6. Click on the **Define Additional Properties** option.
7. Click on the **Create** button.
8. Click on the **Members** property page button.
9. Click on the **Add** button.
10. Select a desired Distribution List member object from the Objects list box and click on **OK**.

    Repeat Steps 9 and 10 until all the desired objects have been selected.

11. Click on **OK**.

External Entity

The External Entity object represents non-native NDS objects. This object is an NDS place holder that allows you to send messages to users who would otherwise not be listed in NDS.

Since External Entity objects are normally imported into NDS for use with gateway software, you will not need to use External Entity objects in a basic MHS Services installation.

User, Group, Organizational Role, and Organizational Unit

The User, Group, Organizational Role, and Organizational Unit objects are used for various purposes within NDS. These objects must be assigned mailboxes. You can use either of the following procedures to assign mailboxes to these objects:

Procedure Number 1

1. Using NetWare Administrator, double-click on the **Messaging Server** icon.
2. Click on the **Users** property page button.
3. Click on the **Add** button.
4. Select a desired object from the Objects list box and click on **OK**.

   Repeat Steps 3 and 4 until all the desired objects have been selected.
5. Click on **OK**.

Procedure Number 2

1. Using NetWare Administrator, open or create a User, Group, Organizational Role, or Organizational Unit object.
2. Click on the **Define Additional Properties** option.
3. Click on the **Create** button.
4. Click on the **Mailbox** property page button.
5. Click on the **Browse** button to the right of the Mailbox Location field.
6. Select the appropriate Messaging Server object from the Objects list box and click on **OK**.
7. Click on **OK**.

Please note that all the procedures that use NetWare Administrator may also be performed using NETADMIN. Only one utility has been documented here.

Using FirstMail to Send and Receive E-Mail

MHS Services for NetWare 4 includes MS Windows and DOS versions of Novell FirstMail. These starter E-mail applications are automatically installed in the SYS:\PUBLIC directory of your messaging server. You do not need to configure FirstMail since it is NDS-aware. As soon as MHS Services is installed and running, your network users can begin using messaging services.

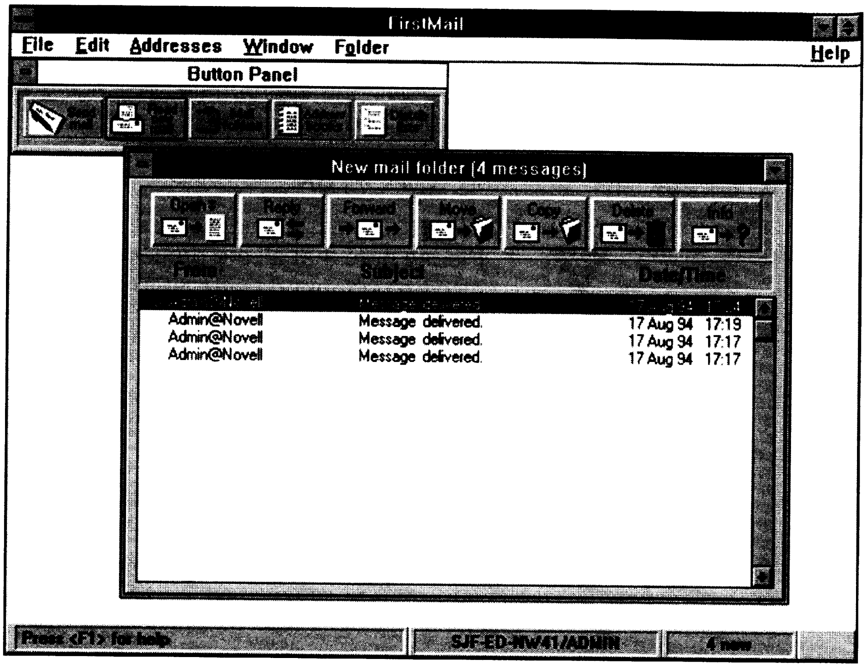

From	Subject	Date/Time
[illegible]	[illegible]	[illegible]
Admin@Novell	Message delivered.	17 Aug 94 17:19
Admin@Novell	Message delivered.	17 Aug 94 17:17
Admin@Novell	Message delivered.	17 Aug 94 17:17

Figure 8-7: FirstMail

Both DOS and MS Windows versions of FirstMail are highly intuitive. General instructions have been provided in the *MHS Services* for *NetWare 4* manual; no further information is provided here.

If you are using a mail-enabled application that is not NDS-aware, you must register the application with MHS Services for NetWare 4 and add the application name to each user's list of applications.

For more information on FirstMail, see "Getting Started with FirstMail" and "Assigning Applications to Users" in *MHS Services for NetWare 4.*

Exercise 8-1: Identifying and Describing the MHS Installation

Activity: Review key information about the MHS Services installation.

Procedure: Answer the following questions. To answer questions 5 through 7, use NetWare Administrator to browse the Directory tree and the file system.

1. What is the command that invokes the installation process?

2. How does installing MHS Services after installing NetWare 4 differ from installing MHS Services as a part of the initial NetWare 4 installation?

3. During the installation, you are prompted to enter the name of a user that administers MHS; what is this user called?

4. Is the name for the user who administers MHS the user's name or the user's full distinguished name? Example: Admin or Admin.CORP.

5. What NDS objects are created during the MHS Services installation?

6. What subdirectory is FirstMail for Windows and FirstMail for DOS stored in?

7. What is the default path for mailbox assignments to users?

8. What is the command you can add to a file to invoke MHS when the server is started?

9. What is the name of the file to which the command from the previous question is added?

Exercise 8-2: Assigning Mailboxes to NDS Objects

Activity: Using NetWare Administrator, assign mailboxes to User, Distribution List, and Organizational Role objects. Create a Distribution List.

Procedure: Use NetWare Administrator to perform the following:

1. Log in as Admin.EMA*nnn*.
2. Assign a mailbox to a User object using the Messaging Server object.
   a. Double-click on the **Messaging Server** object.
   b. Click on the **Users** property page button.
   c. Click on the **Add** button.
   d. Select a desired object from the Objects list box and click on **OK**.
   e. Repeat Steps 2c and 2d until all the desired objects have been selected.
   f. Click on **OK**.
3. Create a new user object and assign a mailbox to the user.
   a. Click on the **EMA*nnn*** container object.
   b. Access **Create** from the Objects menu.
   c. Select the **User** icon and click on **OK**.
   d. Click on the **Define Additional Properties** option.
   e. Fill in the **Login Name** and **Last Name** fields.
   f. Click on the **Create** button.
   g. Click on the **Mailbox** property page button.
   h. Click on the **Browse** button to the right of the Mailbox Location field.
   i. Select the appropriate Messaging Server Object from the Objects list box and click on **OK**.
   j. Click on **OK**.

4. Create a new Organizational Role called **MHS**.

5. Make your assigned User object the Occupant of the Organizational Role.

   a. Click on the Organizational Role object you created and access **Details** from the Object menu.

   b. Click on the **Browse** button to the side of the Occupant field.

   c. Click on the **Add** button.

   d. Click on your assigned User object from the Objects list.

   e. Click on **OK**.

6. Assign a Mailbox property to the EMA*nnn*_MSG Messaging Server object.

   a. Click on the **Mailbox** page button.

   b. Click on the **Browse** button to the side of the Mailbox Location field.

   c. Highlight the **EMA*nnn*_MSG** Messaging Server object in the Objects list.

   d. Click on **OK**.

7. Create a Distribution List object and name it **DLIST**.

8. Assign two User objects, including the User object you created in Step 3, to the Distributed List.

   a. Highlight the **DLIST** object you created and access Details from the object menu.

   b. Click on **Members**.

   c. Click on **ADD.**

   d. Select the **User** object from the Objects list.

   e. Click on **OK**.

   f. Repeat Steps c, d, and e to add the other User object.

Exercise 8-3: Sending and Receiving E-Mail Using FirstMail

Activity: Create, send, and read messages using FirstMail.

Procedure: Using FirstMail for Windows, perform the following tasks:

1. In MS Windows, double-click on the **FirstMail** icon to launch the program.
2. Create a message and address it to the User object that you created in the previous exercise.
   a. Select **New Message** from the File menu or click on the first icon (pen and paper) on the button bar.
   b. In the To: field, type the User object name.
   c. In the Subj: field, type a title to your message.
   d. Click in the message area and type a message of your choice.
   e. When you have completed your message, click on the **Send** button.
3. Create a message and address it to a Distribution List object.
   a. Select **New Message** from the File menu or click on the first icon (pen and paper) on the button bar.
   b. In the To: field, type the Distribution List object name. Example: To: **DLIST**.
   c. In the Subj: field, type a title to your message.
   d. Click in the message area and type a message of your choice.
   e. When you have completed your message, click on the **Send** button.
4. Exit FirstMail and MS Windows.
5. Log in as your assigned User object.
6. Start MS Windows and launch FirstMail for Windows.

7. Access **Read New Mail** from the File menu or click on the second icon (envelope) on the button bar.
8. Double-click on the message to read it.
9. Save the message.
   a. Make sure the message is highlighted.
   b. Click on the **Move** button in the "New mail folder" button bar.
   c. Select the **Mail Folder**. Your message will be stored in this folder.

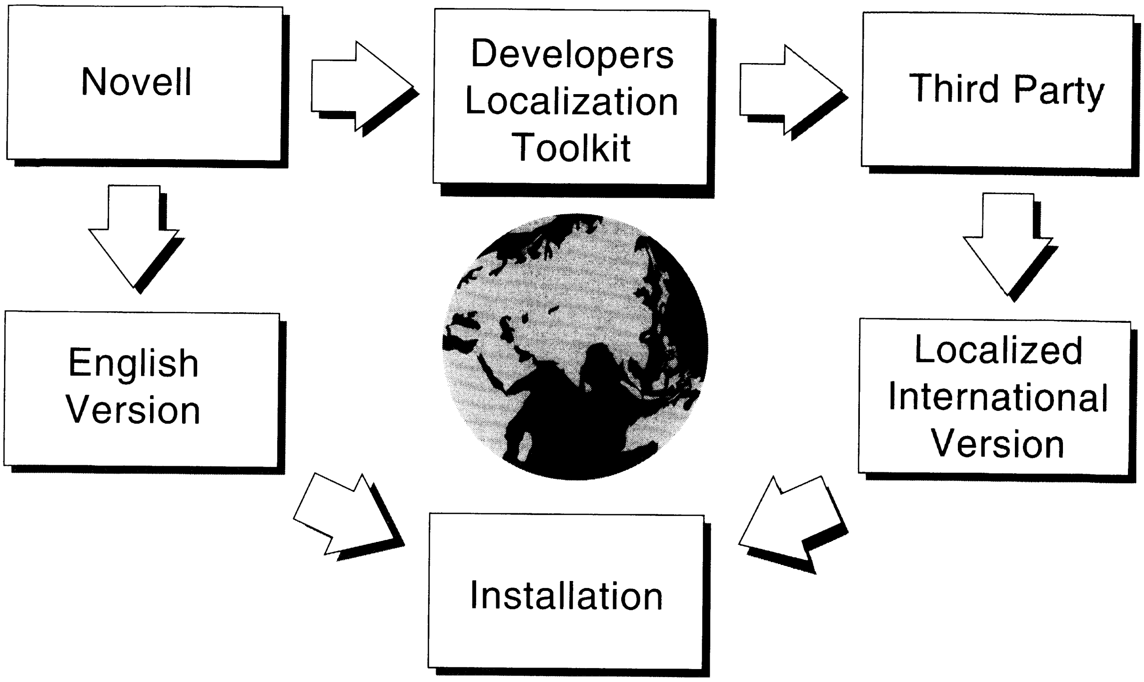
Novell
Developers Localization Toolkit
Third Party
English Version
Localized International Version
Installation

Figure 8-8

Internationalization

The internationalization feature makes NetWare 4 easy to use with multiple languages and cultures. Language support, or internationalization, includes the following:

- Text displayed in non-English languages for utilities and system messages
- Alternate display formats for numbers, dates, times, and other parameters that appear differently in different countries
- Double-byte character support for Asian languages (part of the Unicode specification)
- Non-English filename characters and path separators

Language Enabling

Language enabling means that NetWare 4 is structured to isolate all message strings (such as error messages and menu items) from the source code and place them in separate files called Language Modules. If message strings were left in the source code, developers would need to recompile and relink the source code to change the language of the messages.

All message strings in NetWare 4 source code are replaced with pointers to the location of the strings in the Language Module. Language Module filenames are the names of the programs plus .MSG as the filename extension. For example, the Language Module for the MONITOR.NLM is MONITOR.MSG. Language Modules for different languages all have the same name; therefore, they must be kept in separate directories.

Setting the Language for the Server

To specify the language of the NetWare operating system, do the following:

1. Bring down the server.
2. Place the SERVER.MSG file for the desired languages in the NetWare server boot directory.
3. Run SERVER.EXE.

Changing an NLM Language

Changing the language of the server will change the language for subsequently loaded NLMs. The language for the NLMs is determined by the language designator, which directs the NLM loader to the NLS subdirectory for the language module.

To change the language, you use the LANGUAGE console command. To use LANGUAGE, type the following:

LANGUAGE *language designator* <Enter>

The figure below lists a number of languages and their corresponding language designators.

0	Canadian French	**8**	Italian
1	Chinese	**9**	Japanese
2	Danish	**10**	Korean
3	Dutch	**11**	Norwegian
4	English	**12**	Portuguese
5	Finnish	**13**	Russian
6	French	**14**	Spanish
7	German	**15**	Swedish

Figure 8-9: Language Designators

Changing the NetWare Server Keyboard Type

NetWare 4 also allows you to use server keyboard types other than United States English by loading the KEYB.NLM. The following types are included in the initial release of NetWare 4:

- United States (English)
- Germany
- France
- Italy
- Spain

To view a list of valid keyboard settings, type the following:

LOAD KEYB <Enter>

The keyboard types are dependent on the code page set by DOS. Most keyboard types work with the default code page. If you must change your DOS code page setting, see your DOS manual for more information.

NLM Message Search Hierarchy

The NLM search hierarchy is based on the fact that every NetWare 4 server has a default language for loading NLM messages and help files. This language is based on the language designator.

The NLM search hierarchy uses the language designator in connection with the directory structure shown in Figure 8-10. The numbered subdirectories under the NLS directory correspond to the NetWare server language designators. For example, Danish language modules are stored in directory 2.

This directory structure allows you to load various languages on a NetWare 4 server. The INSTALL program automatically puts language modules in the appropriate NLS subdirectories each time a new international version is installed.

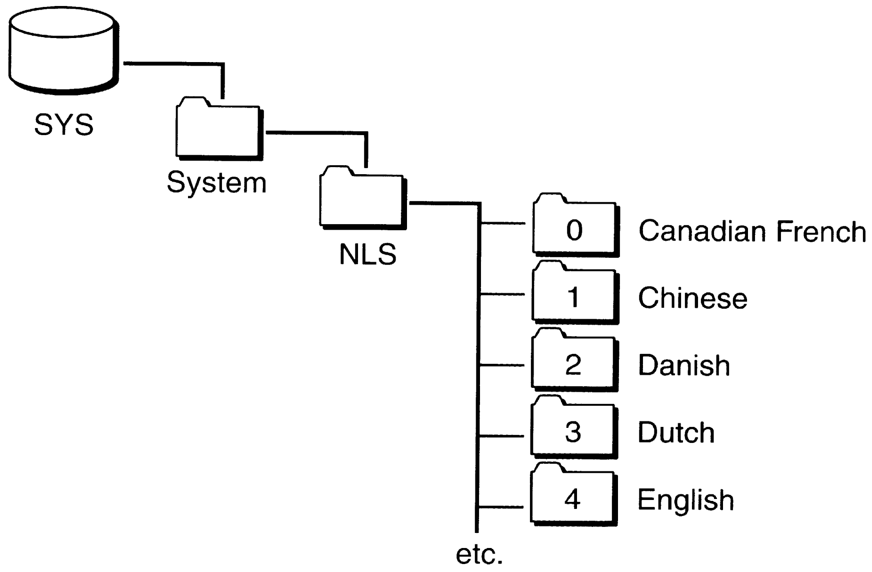

Figure 8-10: Search Hierarchy

Configuring a Server for Locale Formats

The formats for expressing time, dates, and numbers differ among locales. Locale format information is stored in LCONFIG.SYS. Each international version is shipped with an LCONFIG.SYS file. The LCONFIG.SYS file is found in the C:\NWSERVER directory; it is set to the formats of the locale for that international version.

Locale	Date	Time	Number
USA	09/19/90	8:37:00 PM	12,345.67
UK	19/09/90	20:37:00	12,345.67
France	19.09.90	20:37:00	12 345,67
Germany	19.09.90	20:37:00	12.345,67

Figure 8-11: Locale Formats

Configuring a Workstation for Languages

To configure a client workstation for a non-English language, type the following at the workstation:

SET NWLANGUAGE = *[language]*

Replace the variable *[language]* with the name of the language. Languages include English, Duetsch, Espanol, Francais, and Italiano.

You must set this variable to access online help within utilities.

You can automate this process by placing the command statement in the AUTOEXEC.BAT file.

NWLANGUAGE is a DOS environment variable that tells the client to look for a specific message file for the utilities. If no message files exist, the default English message files are used. This variable is also used by DynaText in accessing the proper online manuals.

Summary

NetWare 4 includes tools that allow you to communicate with dissimilar types of computer systems over wide distances. The NetWare MultiProtocol Router manages wide-area communications between various types of systems.

MHS Services for NetWare 4 provides a full-featured mail handling system that allows storage, forwarding, and delivery of electronic mail from many types of mail systems. NetWare 4 includes FirstMail.

NetWare 4 provides language-specific modules that allow users to view utilities, server messages, and online documentation in their native language.

SECTION 9 Optimizing the Network and Server

Introduction

Your responsibilities as network administrator go beyond day-to-day management; you are also responsible for ensuring that your network performs at optimum efficiency.

To help you achieve this goal, NetWare 4 tracks key performance indicators and provides you with the related performance statistics.

This section explains concepts related to the performance statistics and gives you the information you need to interpret the statistics and take action on them. It also describes features designed to improve performance and factors that influence performance.

In addition, this section explains parameters that you can modify to improve performance, and provides procedures and suggestions that help you optimize performance on your network.

Objectives

At the end of this section, you will be able to do the following:

1. Describe NetWare 4 memory management architecture, including memory allocation and memory protection features.
2. Interpret the MONITOR Statistics screen.
3. Monitor and modify file and directory cache performance.
4. View and modify server buffer and packet parameters.
5. Define and enable memory suballocation.
6. List the steps to enable file compression.
7. Enable and manage the Packet Burst protocol.
8. Enable and manage Large Internet Packets (LIPs).

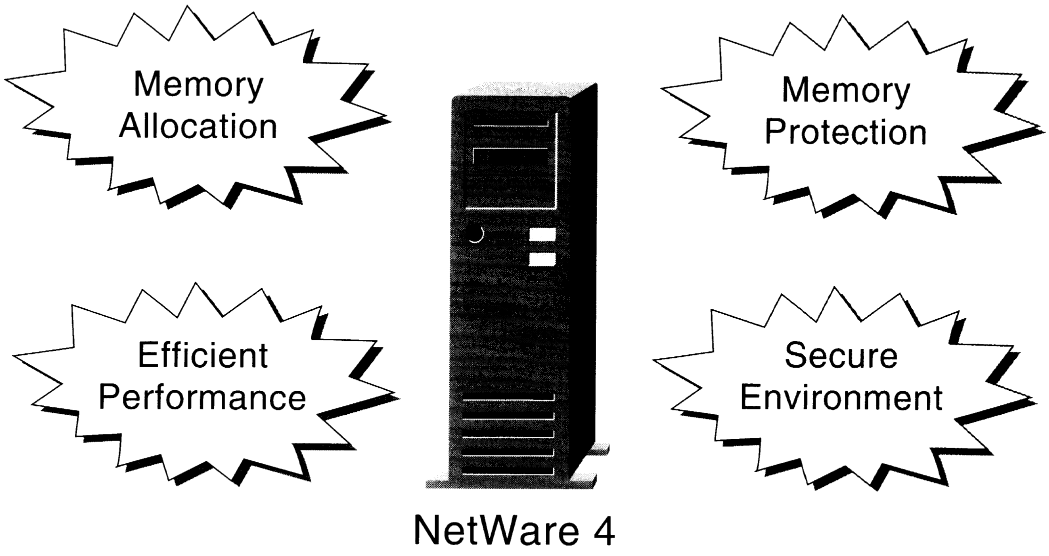

Figure 9-1

NetWare 4 Memory Management

As network administrator, you need to be familiar with memory allocation and memory protection and their implications for memory management. Most memory management is handled by the NetWare 4 memory management architecture. Memory allocation and memory protection features provide efficient performance and provide a simple and straightforward environment for NLM developers. If you need to fine-tune your system, you can use MONITOR, SET, and SERVMAN.

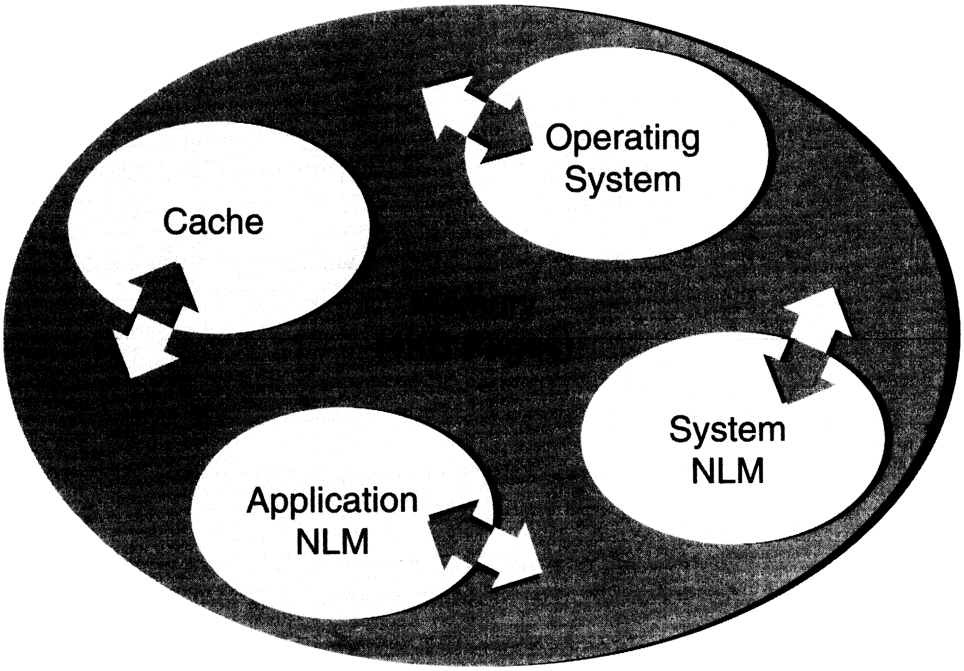

Figure 9-2

Memory Allocation

When you manage the network, you need to be aware of memory allocation and how it affects NetWare server performance.

NetWare uses paging to allocate all memory resources in the system. A page is a 4KB block of RAM. The system can assign physically discontiguous pages of memory in a logically contiguous (adjacent) range, resulting in efficient memory utilization.

Memory allocation reserves a specific memory location in RAM for processes, instructions, and data. In the NetWare 4 memory architecture, the network operating system (NOS) gives each process its own allocation pool. Each process requests memory from its allocation pools and returns memory to the same pool; this minimizes fragmentation and maximizes efficiency. Memory assigned to process allocation pools is mapped to the global pool of memory pages.

NetWare 4 memory allocation also allows the expedient development of third-party applications and the optimization of NLMs. The combined function calls (APIs) in the memory management scheme are optimized for NLMs because NLMs tend to make a set number of allocations and use that pool of memory over and over during the life of the process.

Memory Deallocation and Garbage Collection

NetWare 4 uses memory deallocation and garbage collection to collect unused segments of memory and return them to a common memory pool. Reclaiming unused memory nodes is important to maintaining performance.

First, memory must be deallocated by the API named Free. This process simply labels the memory as deallocated.

Then the deallocated memory must be recovered. Garbage collection recovers the pieces of deallocated and available system memory so system memory does not become depleted over time.

The garbage collection routine is interruptible, can run in the background, and should be run frequently.

You can use three server console SET commands to optimize garbage collection. Use these commands to perform the following tasks:

- Garbage Collection Interval – sets the interval for collection. The range is from 1 to 60 minutes; the default is 15.
- Number of Frees for Garbage Collection – sets the number of calls by each NLM that triggers garbage collection. The range is 100 to 100,000. The default is 5,000.
- Minimum Free Memory for Garbage Collection – sets the minimum number of bytes that must be available in the alloc memory pool for successful garbage collection. The range is from 1,000 to 1,000,000. The default is 8,000.

For further information about freeing memory, refer to *Supervising the Network.*

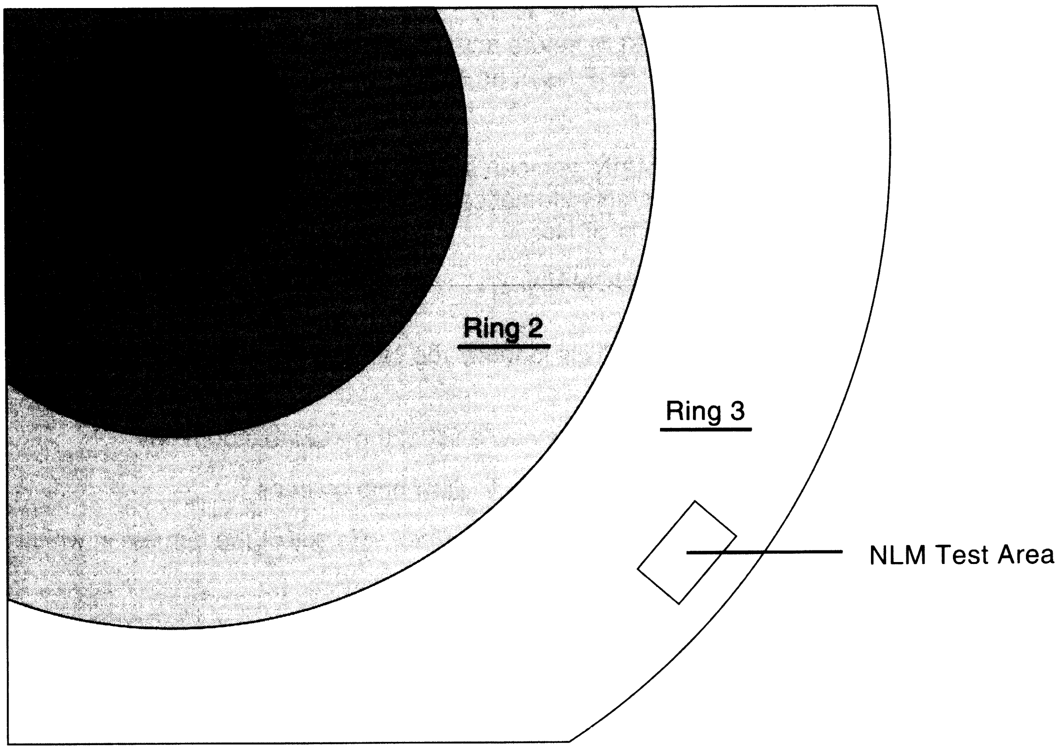

Figure 9-3

Memory Protection

In NetWare 4, memory resources are structured so that processes are kept from corrupting each other's memory space or that of the NOS. NetWare 4 also allows you to establish a safe area or *domain* (a legal addressing space for a process) in which to run untested third-party NLMs until you are confident of the NLMs' integrity. This domain is called the Operating System Protected (OSP) Domain. Developers can debug and test NLMs in this domain before running the NLMs in an unprotected production environment. NLMs loaded into the OSP cannot corrupt the NOS in its OS Domain.

Enabling Memory Protection

To create the OSP and enable memory protection, load the DOMAIN.NLM before any other NLMs by including it in the STARTUP.NCF file. If an NLM fails, it may not be possible to unload the domain.

With DOMAIN, you can specify the ring you will use to load the NLM. To load DOMAIN automatically, include the following line in the STARTUP.NCF file:

```
LOAD DOMAIN
```

Once DOMAIN is loaded, you can use the following DOMAIN parameters:

- DOMAIN – To view a list of the available domains
- DOMAIN HELP – To view help screens
- DOMAIN = <*domainname*> – To select the domain in which you will load NLMs

Existing NLMs may need to be modified to function properly within memory protection.

Ring Protection

An important facet of memory protection is protection between rings. The ring protection model is a logical model of concentric circles that illustrates how programs can run at different privilege levels, or rings, remaining isolated from one another. This isolation prevents the processes from inadvertently accessing, and possibly conflicting with, NLMs in the OS Domain.

Each ring has a different privilege level. The OS Domain runs in ring 0, where it has the highest privilege. Ring 0 does not have segment limits and read/write restrictions; therefore, code and data are not protected. The OSP Domain runs in ring 3. A process running in one ring cannot gain access to memory at a greater privilege level (lower numerical level) unless explicitly allowed.

Once an NLM has been tested extensively and appears stable, the NLM can be moved to ring 0, the OS Domain, and run from there for a gain in performance. At ring 0, NLMs generally run faster because the processes are trusted and are not slowed down by validation processes.

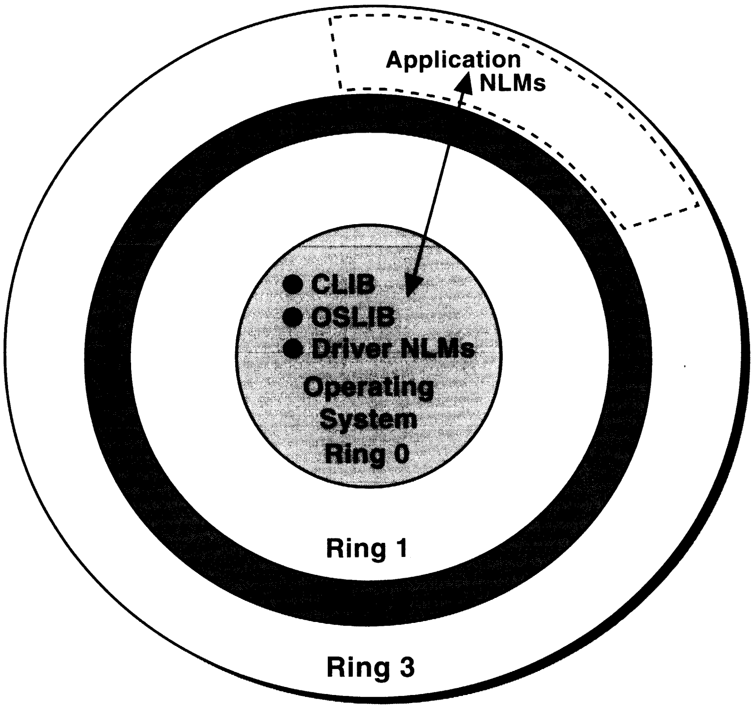

Figure 9-4: Ring Protection Model

Segmentation

Individual processes within a ring are protected from one another through segmentation. Segmentation limits the memory a process can address. These size limits, implemented as a set of bits in the code and data segments, create a fence around the system's domain. This fence prohibits the processor from executing instructions or accessing data above a certain memory location for each process.

Benefits of Memory Protection

The memory protection scheme used in NetWare 4 provides flexible, yet full, protection for the operating system and all NLMs. Network administrators and third-party developers can implement the protection mechanism according to their needs.

Memory protection in NetWare 4 provides the following benefits:

- Allows testing of new services and protects against loss of data due to ill-behaved NLMs
- Secures services and operating system environments
- Allows graceful termination for services that fail
- Allows services in ring 0 to keep running during failure of protected domain NLMs

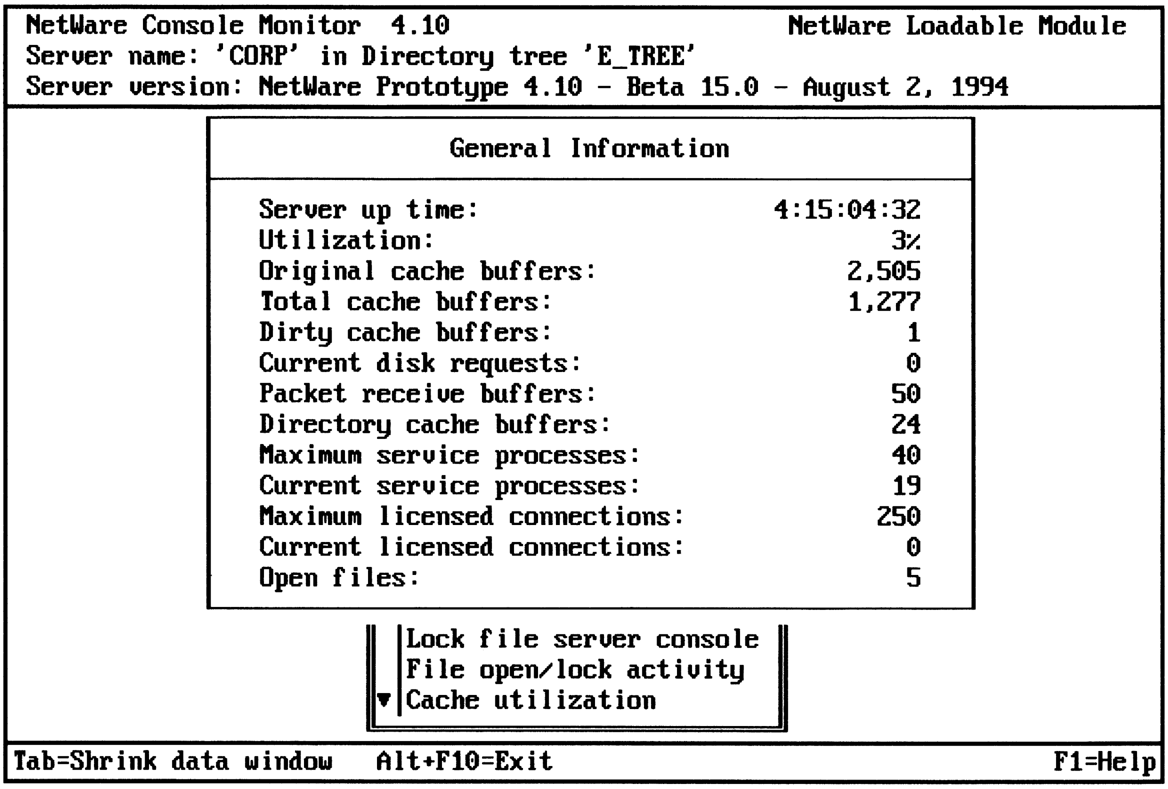

```
NetWare Console Monitor  4.10                    NetWare Loadable Module
Server name: 'CORP' in Directory tree 'E_TREE'
Server version: NetWare Prototype 4.10 - Beta 15.0 - August 2, 1994

                        General Information

            Server up time:                          4:15:04:32
            Utilization:                                     3%
            Original cache buffers:                       2,505
            Total cache buffers:                          1,277
            Dirty cache buffers:                              1
            Current disk requests:                            0
            Packet receive buffers:                          50
            Directory cache buffers:                         24
            Maximum service processes:                       40
            Current service processes:                       19
            Maximum licensed connections:                   250
            Current licensed connections:                     0
            Open files:                                       5

                       Lock file server console
                       File open/lock activity
                      ▼Cache utilization

Tab=Shrink data window   Alt+F10=Exit                           F1=Help
```

Figure 9-5

MONITOR Statistics

The main MONITOR screen, General Information, displays information about the status of the network. This screen appears for a few seconds after you load MONITOR.NLM. Often you can track the performance of the network from the main MONITOR screen without having to look at specific displays.

The following table lists the statistics available on the MONITOR General Information screen.

Statistic	Explanation
Operating system version and date	Version and release date of the operating system (upper left corner of the screen).
Information for server	Name of the NetWare server and current Directory tree.
Server up time	Length of time the NetWare server has been running since it was last booted.
Utilization	Percentage of time the processor is being used.
Original cache buffers	Number of cache buffers available when the server is first booted. The figure represents the number of blocks installed as memory in your server less the OS kernel and DOS.
Total cache buffers	Number of buffers currently available for file caching. This number decreases as modules are loaded into memory.
Dirty cache buffers	Number of cache buffers containing information that needs to be written to disk.
Current disk requests	Number of disk requests in a queue that the server is waiting to service.
Packet receive buffers	Number of buffers available to receive station requests.
Directory cache buffers	Number of buffers allocated for directory caching.
Service processes (maximum and current)	Number of task handlers allocated for station requests. When the number of station requests exceeds what is currently allocated, the operating system employs extra task handlers to execute the requests. Once memory is allocated for a service process, it remains allocated. The only way to remove this memory allocation is to bring down the server.
Connections In Use (maximum and licensed)	Number of stations supported by the current license and the number currently attached to the NetWare server.
Open files	Number of files being accessed by the NetWare server and by network stations.

Table 9-1: MONITOR Statistics

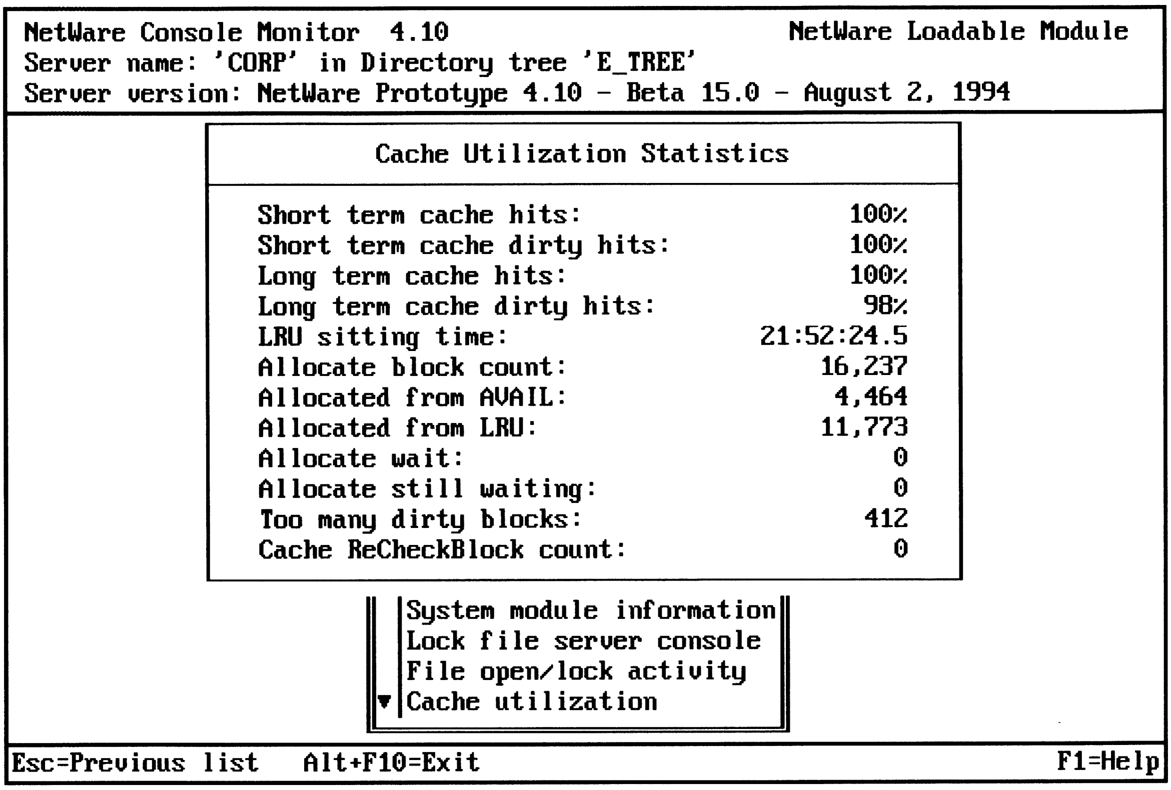

```
NetWare Console Monitor  4.10                    NetWare Loadable Module
Server name: 'CORP' in Directory tree 'E_TREE'
Server version: NetWare Prototype 4.10 - Beta 15.0 - August 2, 1994

                     Cache Utilization Statistics

           Short term cache hits:                        100%
           Short term cache dirty hits:                  100%
           Long term cache hits:                         100%
           Long term cache dirty hits:                    98%
           LRU sitting time:                       21:52:24.5
           Allocate block count:                       16,237
           Allocated from AVAIL:                        4,464
           Allocated from LRU:                         11,773
           Allocate wait:                                   0
           Allocate still waiting:                          0
           Too many dirty blocks:                         412
           Cache ReCheckBlock count:                        0

                           System module information
                           Lock file server console
                           File open/lock activity
                         ▼ Cache utilization

Esc=Previous list   Alt+F10=Exit                                     F1=Help
```

Figure 9-6

Cache Utilization

NetWare uses caching to speed up reads and writes on server hardware. A cache is a temporary storage area that holds blocks of disk space in RAM. As long as a block of disk space stays in RAM, NetWare can access it much faster than accessing the disk.

A cache hit means that the server found the data in RAM, rather than having to go to the disk. This represents a large performance increase. Cache utilization shows the percentage of cache hits.

Having sufficient RAM to cache disk data results in a high percentage of cache hits. The more NLMs and other services you load, the less RAM is available for caching disk blocks. This lowers your cache hit percentage and slows network performance.

To view Long Term Cache Hits, select **Cache Utilization** from the MONITOR Available Options menu. If the Long Term Cache Hits statistic falls below 90 percent, you should add more RAM. Short-term solutions include using the REMOVE DOS command and unloading NLMs that are not critical at the time.

Cache Buffers and Volume Block Size

You can increase cache utilization and improve response time by modifying cache buffer and volume block size.

The system is optimized when the cache buffer size and volume block size are the same. This may not be possible because NetWare requires that the cache buffer size be the same as the smallest configured volume block size, and each volume on a server may be configured to use a different block size.

Setting Cache Read and Write Parameters

You can improve server performance by increasing the speed of reads and writes. You can adjust the performance of both read and write operations by following the guidelines in the tables below. The guidelines may differ slightly from system to system.

Setting the Write Cache

If network users frequently make many small write requests, set the following:

Parameter	Setting
Dirty Disk Cache Delay Time	7
Maximum Concurrent Directory Cache Writes	25
Dirty Directory Cache Delay Time	2
Maximum Concurrent Disk Cache Writes	50

Setting the Read Cache

If your server is slow to respond to read requests, set the following:

Parameter	Setting
Maximum Concurrent Disk Cache Writes	10
Maximum Concurrent Directory Cache Writes	5
Directory Cache Buffer NonReferenced Delay	60

These parameters can be changed using either SERVMAN or SET.

Setting read and write cache performance levels may offset each other. If write cache is set too high, read cache may be affected.

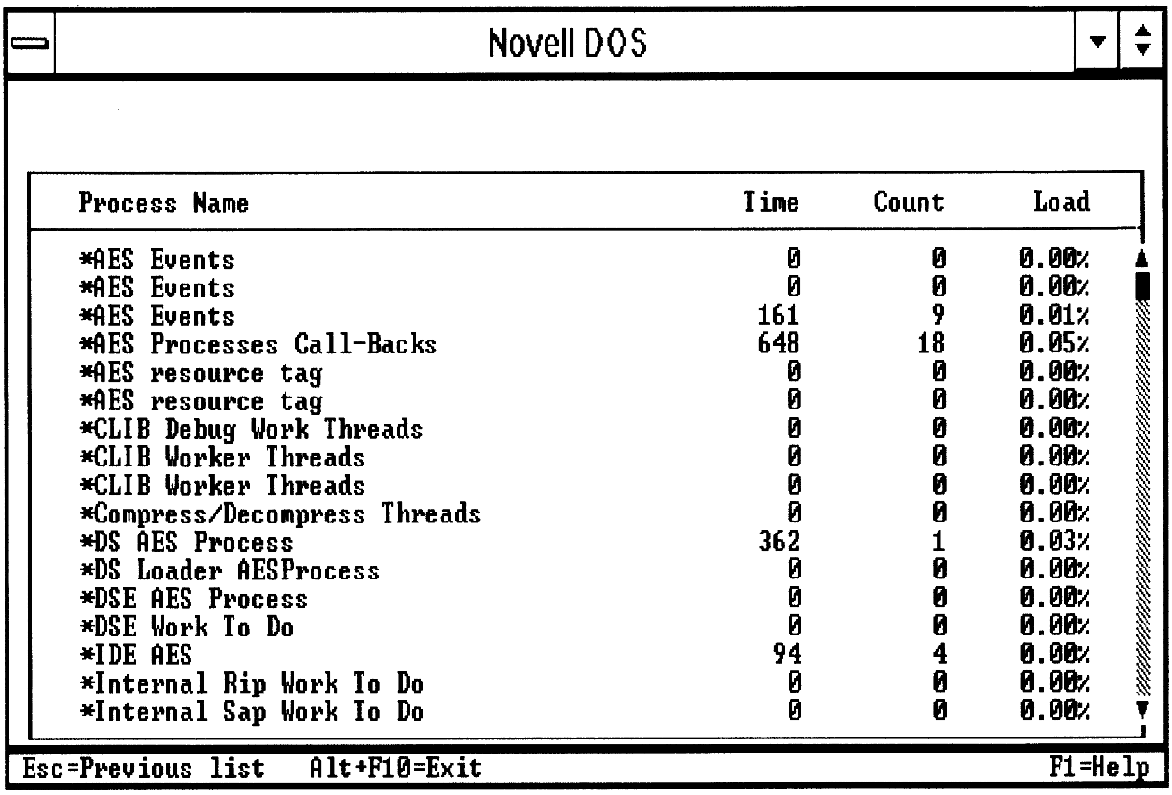

```
Process Name                           Time      Count      Load

*AES Events                               0          0     0.00%
*AES Events                               0          0     0.00%
*AES Events                             161          9     0.01%
*AES Processes Call-Backs               648         18     0.05%
*AES resource tag                         0          0     0.00%
*AES resource tag                         0          0     0.00%
*CLIB Debug Work Threads                  0          0     0.00%
*CLIB Worker Threads                      0          0     0.00%
*CLIB Worker Threads                      0          0     0.00%
*Compress/Decompress Threads              0          0     0.00%
*DS AES Process                         362          1     0.03%
*DS Loader AESProcess                     0          0     0.00%
*DSE AES Process                          0          0     0.00%
*DSE Work To Do                           0          0     0.00%
*IDE AES                                 94          4     0.00%
*Internal Rip Work To Do                  0          0     0.00%
*Internal Sap Work To Do                  0          0     0.00%
Esc=Previous list   Alt+F10=Exit                              F1=Help
```

Figure 9-7

Processor Utilization

Another critical factor in network performance is the CPU. A heavy load on the CPU caused by one or more processes can dramatically impact network performance.

To view a utilization histogram, select **Processor Utilization** from the MONITOR Available Options menu; then press <F3>. Note in the Load column which processes are using the most CPU time.

The Idle Loop process also shows processor utilization. Whenever the CPU is fairly idle, the idle loop should register percentages over 90%.

The percentage of idle loops is inversely related to the utilization percentage; when processor utilization is high, idle loop percentages fall.

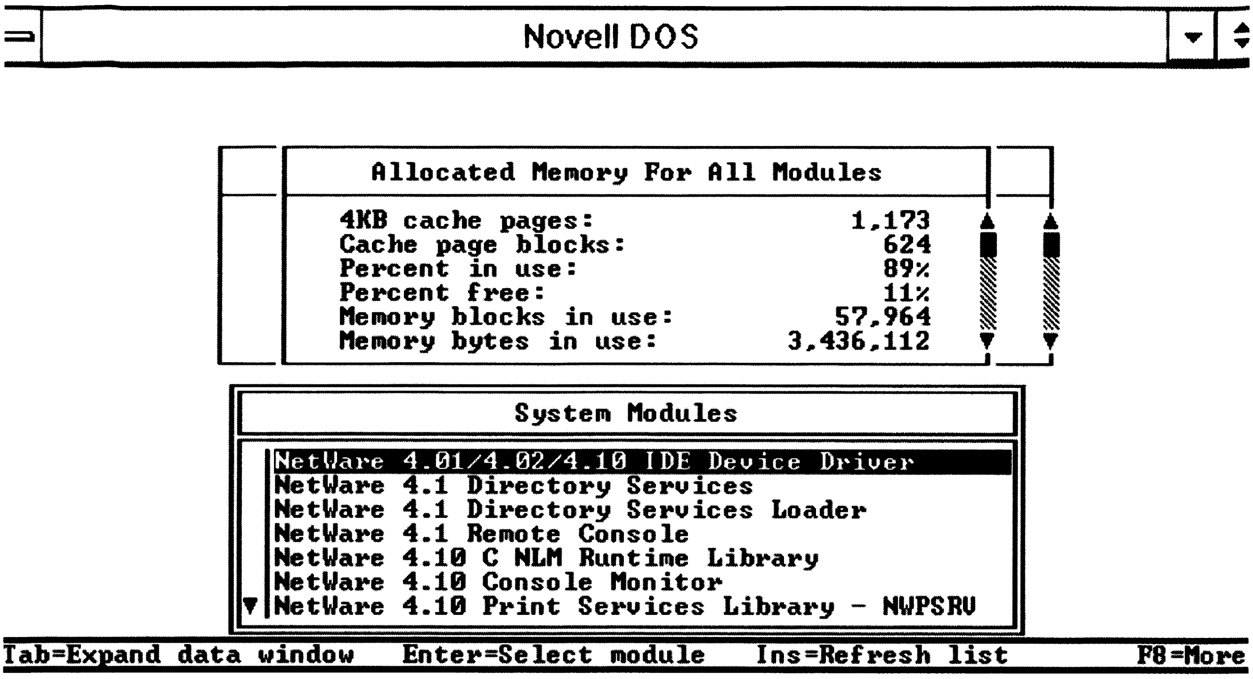

```
Allocated Memory For All Modules

4KB cache pages:                    1,173
Cache page blocks:                    624
Percent in use:                       89%
Percent free:                         11%
Memory blocks in use:              57,964
Memory bytes in use:            3,436,112

System Modules

NetWare 4.01/4.02/4.10 IDE Device Driver
NetWare 4.1 Directory Services
NetWare 4.1 Directory Services Loader
NetWare 4.1 Remote Console
NetWare 4.10 C NLM Runtime Library
NetWare 4.10 Console Monitor
NetWare 4.10 Print Services Library - NWPSRV

Tab=Expand data window   Enter=Select module   Ins=Refresh list        F8=More
```

Figure 9-8

Memory Utilization

Available memory is a factor in performance, since it affects the number of cache buffers. The NLM loading and unloading process provides flexibility in managing memory to ensure high performance. NLMs immediately return memory when they are unloaded.

Keep the following points in mind:

- Many NLMs rely on prerequisite NLMs. Memory needed for loading an NLM is approximately the same as the file size of the NLM, plus any prerequisite NLMs.
- Some NLMs allocate additional memory as they are running.

Use the Memory utilization and System module information options in MONITOR to check memory before and after loading an NLM. The difference indicates the approximate size of the NLM. Test new NLMs and monitor changes in server performance. Intermittent problems may indicate that the NLM "grows" as it runs and then returns memory after completing a task.

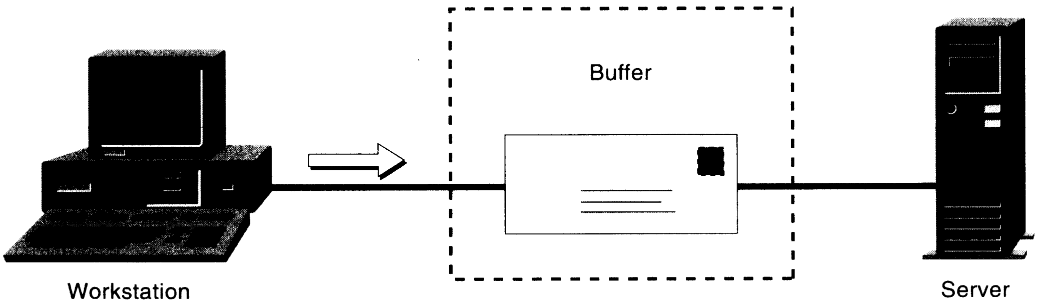

Figure 9-9

Server Packets and Buffers

Server packet and buffer parameters are critical factors. You may need to modify them to enhance server performance.

Packets

Packets are tiny segments of data that are transmitted across a network. Packet size is automatically determined by the NetWare server and the workstation.

The default packet size depends on the capacity of the LAN driver being loaded. The allowable range for a workstation is 576 to 6500 bytes, depending on the LAN driver.

Servers and workstations negotiate the largest packet size available for communications. If the server can handle a 4KB packet and the workstation can handle only a 1KB packet, they will exchange 1KB packets. If the server and workstation are both configured to handle 4KB packets, they can exchange 4KB packets as long as any intermediate routers are also configured to handle 4KB or larger packets.

Maximum Physical Receive Packet Size

Use SERVMAN to modify the Maximum Physical Receive Packet Size parameter to the largest client packet size actually in use on the network. The network board driver should specify the maximum packet size it can handle. This parameter cannot be modified at the console prompt; you must include the parameter in the STARTUP.NCF file.

Ideally, you will use the largest possible packet size, except when the client and server are configured for different packet sizes or when the packet must cross a router to get to the server.

Topology	[illegible]
Token Ring (16 Mbps)	4202
Token Ring (4 Mbps)	4202
Ethernet	1514
ARCnet	4202

Table 9-2: Example Default Maximum Packet Sizes by Topology

Packet Receive Buffers

Packet receive buffers are areas in the NetWare server memory set aside to hold data packets.

Use MONITOR to check current settings. If necessary, you can change the settings using the following parameters.

Maximum Packet Receive Buffers

This parameter determines the maximum number of packet receive buffers that the operating system can allocate.

Often, the default setting (100) is too low for busy servers. A setting between 300 and 500 is ideal for such servers.

Maximum Services Process

A related parameter is Maximum Services Process. Be sure to check this number at the Monitor General Information screen. If the number of allocated service processes is at its maximum, you can increase the maximum to decrease the need for more packet receive buffers.

Minimum Packet Receive Buffers

This parameter determines the minimum number of packet receive buffers that the operating system can allocate. Rather than waiting for requests to come in, the operating system allocates the minimum number of buffers as soon as the NetWare server boots.

If the allocated number is higher than 10 and the server is slow in responding immediately after it has been booted, increase this number. The default is 50.

This parameter must be set in the STARTUP.NCF file.

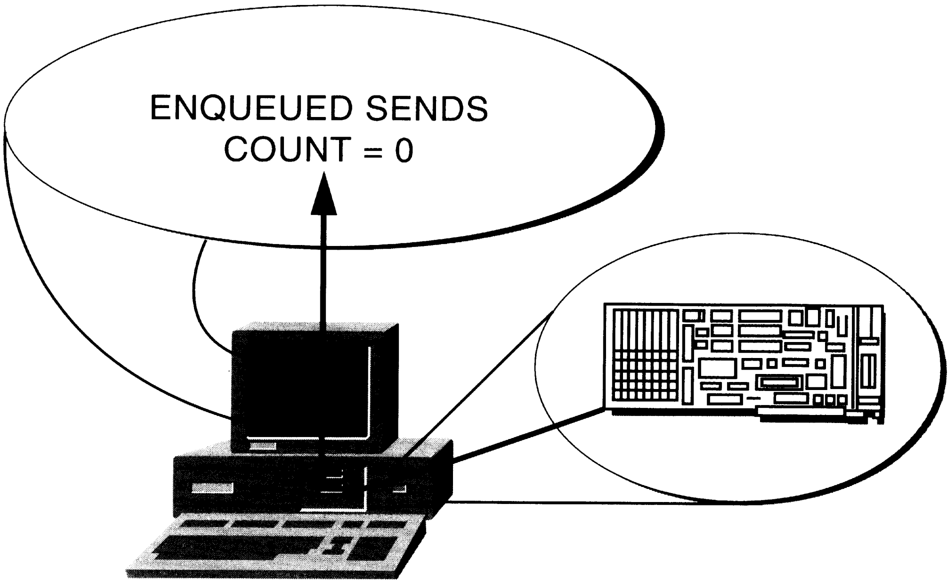

Figure 9-10

Network Interface Board Statistics

The network board handles traffic between the workstation and the server. Network board statistics need to be reviewed to monitor the performance of the network. These statistics reveal trends in communication or problem areas that need to be addressed.

Most network board drivers can provide 15 different statistics, including general error entries, total packets sent and received, and packets dropped due to no available communication buffers. Additional custom statistics may be available, depending on your network board, and may be useful as a housekeeping function.

One helpful custom statistic for the Ethernet NE2000 network board is Enqueued Sends Count. This statistic indicates the number of packets that the board had to buffer because the driver was too busy to send a packet the processor had prepared.

If this count increments regularly, it indicates that the Ethernet NE2000 network board driver is having trouble keeping up with the server and is reaching, or has reached, its saturation point. To correct this, consider replacing an 8-bit or 16-bit network board with a 32-bit board.

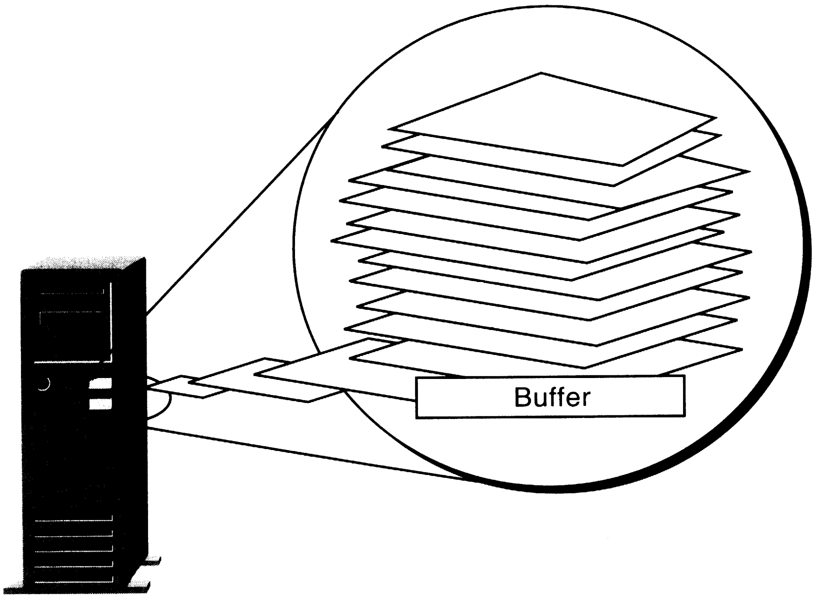

Figure 9-11

Disk Drivers and Controllers

NetWare server performance is affected by the disk drivers and controllers it uses. A disk driver is the interface program that controls hardware. A controller is the hardware that regulates disk drives.

A number of different types of disk drivers and controllers can be used in NetWare servers, as long as they meet Novell standards. High-end servers using the Small Computer Systems Interface (SCSI) boards will use an intelligent device driver, usually supplied by the drive manufacturer. Some of these devices can service multiple drives connected to the same controller.

Some advanced disk controllers employ bus-mastering and 32-bit technologies to optimize the transfer of data to and from the disk.

The main idea to remember is that significant improvements in file read and write times are possible with high-end device driver/controller configurations.

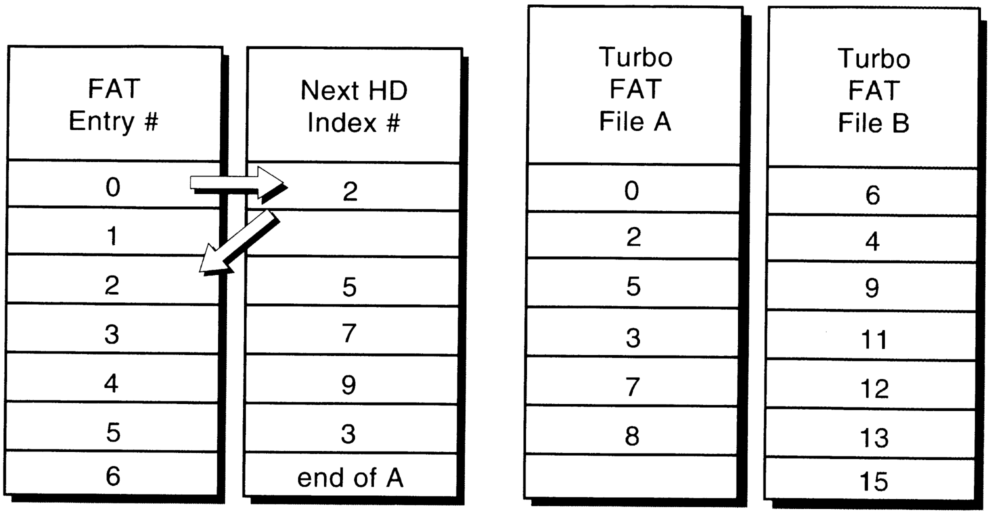

FAT Entry #	Next HD Index #	Turbo FAT File A	Turbo FAT File B
0	2	0	6
1		2	4
2	5	5	9
3	7	3	11
4	9	7	12
5	3	8	13
6	end of A		15

NetWare Server Memory (in RAM)

Figure 9-12

Turbo File Allocation Table (FAT) Indexing

Directory Entry Tables (DETs) and File Allocation Tables (FATs) contain address information that tells the operating system where data can be read from or written to on a volume. Because volumes are divided into disk allocation blocks of a specified size, a file that exceeds the allocation block size is stored in blocks spread over the disk. The system uses a FAT to track the location of files so it can reassemble the files when they are called. A FAT is an index table that points to the disk area where a file is located and links the parts of the file. In NetWare, the FAT is accessed from the DET.

To optimize indexing and accessing very large files, NetWare 4 automatically creates a turbo File Allocation Table (FAT) index for randomly accessed files that exceed 64 FAT entries. Turbo FATs are not created for files that are accessed sequentially, regardless of their size.

The turbo FAT is an index of the blocks and redirections which pertain to that file only. Thus, it does not have to scan through the entire FAT; it can index the table without any scanning, making access to the file faster.

When a file is closed, the turbo FAT index is not flushed immediately from memory. Rather, an aging process begins which involves a timer that determines when a memory space will be flushed and reused. You can alter the Turbo FAT Re-Use Wait Time parameter to help ensure that the index structure will not be flushed if the file is reused quickly.

The default time for the Turbo FAT Re-Use Wait Time parameter is 5 minutes, 29.6 seconds.

- Increase wait time if you know that users frequently reopen the same file after a specific delay and know that other files opened during the delay will reuse the index.
- Decrease the wait time if you want the memory immediately released to service the next file that needs to be indexed.

For more information on FATs, see "File Allocation Table (FAT)" in the *Concepts* manual.

Hands-On Exercise 9-1: Monitoring and Maintaining Network Performance

Scenario

The Manufacturing department is responsible for generating periodic bursts of heavy network traffic. You need to check how this heavy traffic is affecting the users.

Procedure

Use the MONITOR screens to track key performance parameters. Press <F1> at any time for context-sensitive help explaining particular screens and statistics. Write your observations in the space provided. Be prepared to discuss them.

1. Load MONITOR at the server console.

2. Observe the Utilization number in the General Information panel at the top of the main screen. Observe this number as you perform the activities in the following steps.

3. At the server console, note the processor utilization value. Then type **MAP** at the workstation and press <Enter>. On the server console, note the peak processor utilization value.

4. At the server console, note the percentage of Long Term Cache Hits.

5. At the workstation, type

   NDIR *.* /SUB /C <Enter>

6. At the server console, select **Cache Utilization** from the Available Options menu. Note the percentage of Long Term Cache Hits.

7. Press <F1> for an explanation of Short Term Cache Dirty Hits.

8. Select **Processor Utilization** from the Available Options menu and press <F3>. Note which processes use the highest percentage of processor time when users are not active.

9. At the client workstation, create a User object. Note the statistics on the console during the creation of the User object.

10. Select **Scheduling Information** from the Available Options menu and press <Return>. At the client workstation, change to the root of the volume and type **NDIR Z:*.* /SUB**. Compare the changing values of the entries: Interrupts, Idle Processes, and Work.

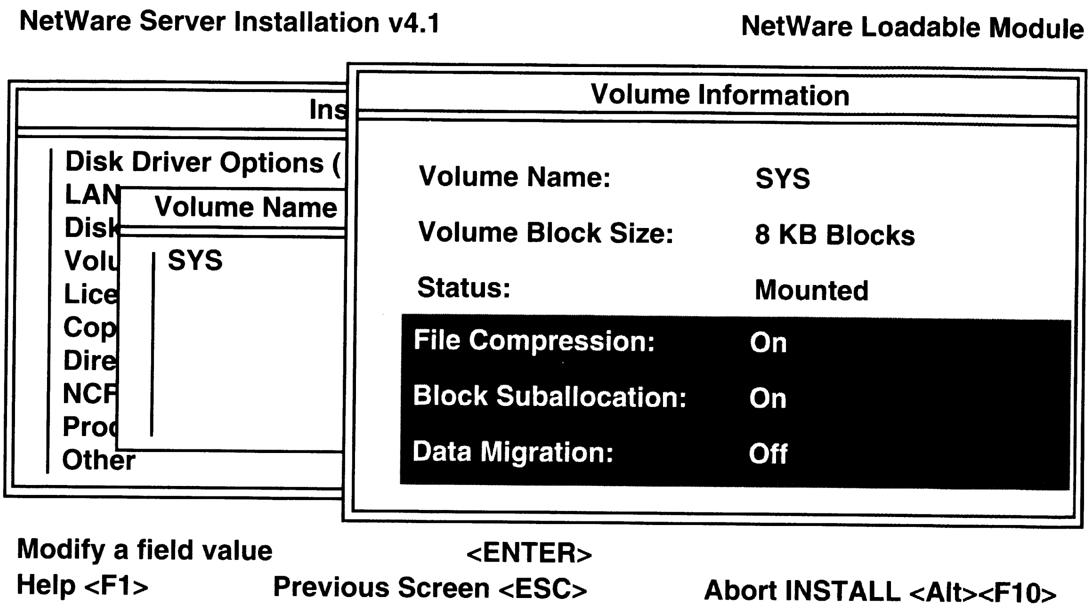

NetWare Server Installation v4.1 NetWare Loadable Module

Ins

Disk Driver Options (
LAN
Disk
Volu
Lice
Cop
Dire
NCF
Proc
Other

Volume Name

SYS

Volume Information

Volume Name:	SYS
Volume Block Size:	8 KB Blocks
Status:	Mounted
File Compression:	On
Block Suballocation:	On
Data Migration:	Off

Modify a field value <ENTER>
Help <F1> Previous Screen <ESC> Abort INSTALL <Alt><F10>

Figure 9-13

Optimizing Disk Space

The memory management schemes discussed so far deal with temporary memory. As a network administrator, you also need to optimize disk space in your file system. NetWare 4 provides the following features to help you optimize permanent memory storage:

- Block suballocation
- File compression

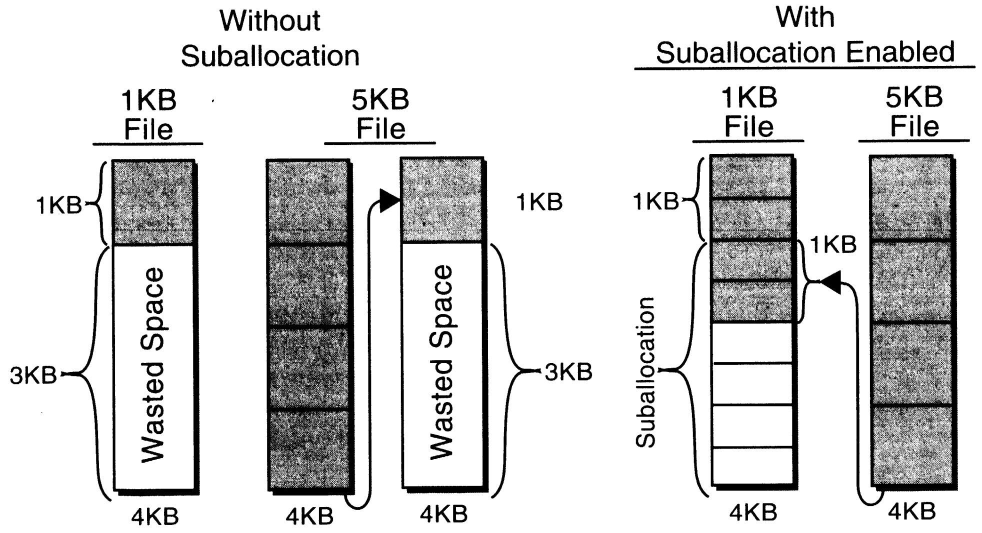

Figure 9-14

Block Suballocation

Block suballocation is a method the system uses to save hard disk space by allowing disk blocks to be allocated in smaller segments.

Files are stored in terms of allocation blocks. Without block suballocation, if a file is smaller than the block size, a full block is still allocated to store the file. For example, a 1KB file on a volume with 4KB allocation block size would still require one 4KB block for storage. The remaining 3KB of that 4KB block are unused and unavailable for storage.

To increase NetWare server speed and ensure more efficient memory usage, NetWare 4 provides 512-byte increments (suballocation blocks) in addition to the volume's allocated block size.

When a file is larger than the allocated block size, the excess is stored in the suballocation blocks (if suballocation is enabled). The beginning of every file is stored at the beginning of a block; if the file exceeds the allocated block size, the remainder of the file is placed in suballocation blocks.

Visualize an allocated block size of 4 KB and a 1KB file. With suballocation, the system stores the 1KB file at the beginning of a 4KB block and uses the remaining 3 KB of that block for suballocation. If a 5KB file is subsequently created, the system stores the first 4 KB of that file in a 4KB block and places the remainder in suballocation blocks.

Block suballocation saves small amounts of disk space per file, which, when multiplied by thousands of files on a large NetWare server, saves large amounts of disk space.

Enabling Block Suballocation

Block suballocation is enabled at the volume level in INSTALL. You cannot enable or disable suballocation whenever you need it; you must enable or disable it at the time you create the volume.

To enable or disable block suballocation, do the following:

1. From the Volume Options, select the volume for which you want to change the block suballocation status.
2. Highlight the **Block Suballocation** field and press <Enter>.
3. Press <Enter> to toggle between **On** and **Off**.
4. Repeat Steps 1 through 3 to modify other volume parameters.
5. Press <Esc> to return to the list of volumes.

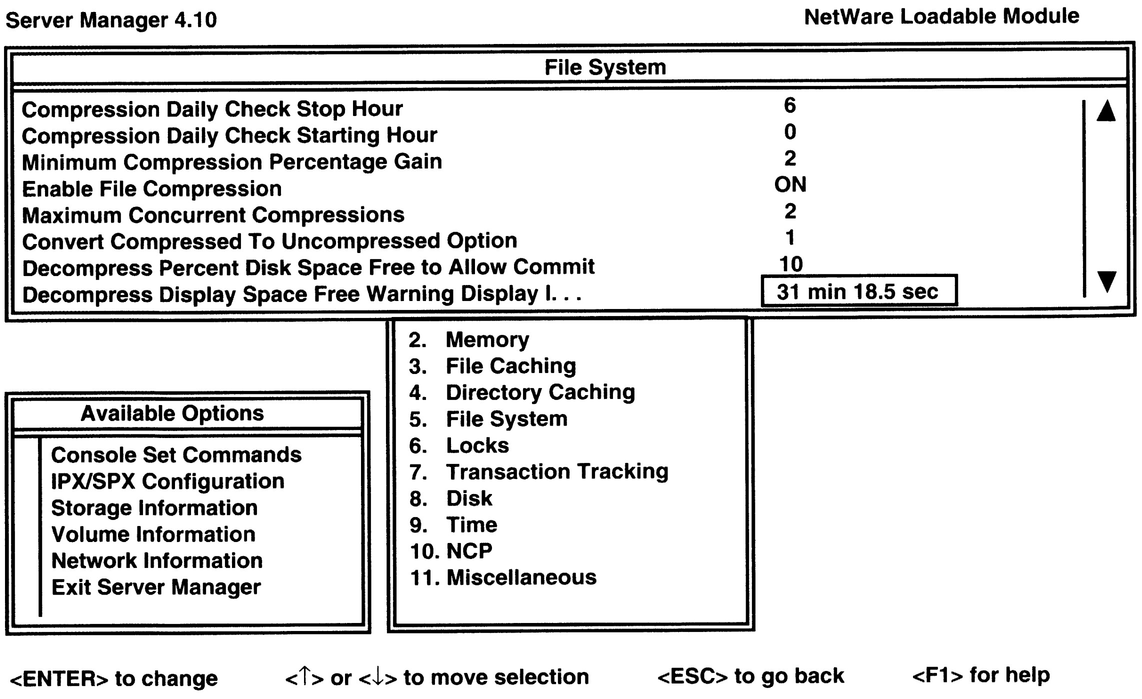

```
Server Manager 4.10                                         NetWare Loadable Module

                                  File System
Compression Daily Check Stop Hour                          6
Compression Daily Check Starting Hour                      0
Minimum Compression Percentage Gain                        2
Enable File Compression                                    ON
Maximum Concurrent Compressions                            2
Convert Compressed To Uncompressed Option                  1
Decompress Percent Disk Space Free to Allow Commit         10
Decompress Display Space Free Warning Display I. . .       31 min 18.5 sec

                                 2.  Memory
                                 3.  File Caching
                                 4.  Directory Caching
     Available Options           5.  File System
 Console Set Commands            6.  Locks
 IPX/SPX Configuration           7.  Transaction Tracking
 Storage Information             8.  Disk
 Volume Information              9.  Time
 Network Information             10. NCP
 Exit Server Manager             11. Miscellaneous

<ENTER> to change    <↑> or <↓> to move selection    <ESC> to go back    <F1> for help
```

Figure 9-15

File Compression

File compression helps you optimize hard disk space by allowing you to store more information on a hard disk. It is a background process that affects NetWare server performance minimally.

File Compression Process

File compression is managed internally by NetWare 4. During file compression, the NetWare 4 operating system does the following:

- Reads and analyzes the file.
- Builds a temporary file describing the original file.
- Determines if any disk sectors can be saved by compressing the file. A gain of at least 2% (default) savings is required before a file is compressed.
- Begins creation of the compressed files.
- Replaces the original files with the compressed files after an error-free compressed version of each file is built.

If a disk error or power failure occurs during compression, the original uncompressed file is retained.

Compression Ratio

The compression ratio is the difference in file size before and after compression. The average compression ratio in a volume is approximately 63 percent. For example, a volume holding 144 MB of data can be compressed to 54 MB.

The file compression/decompression subsystem makes some sacrifice in speed to optimize file compression ratio. It assumes that compression ratio has priority over speed because compression can run during off hours when speed is usually not critical. Rather than compressing every file, the system only compresses files that have not been accessed for some time. Decompression has no noticeable effect on system performance.

Managing File Compression

Managing file compression requires you to perform tasks during installation and after your network has been set up.

During installation, use the INSTALL utility to enable or disable compression per volume and to configure compression for your particular needs. The default is file compression enabled. Once you turn on file compression, you cannot turn it off unless you recreate the volume. However, you can still turn off file compression at the directory and file level; using the SET parameter ENABLE FILE COMPRESSION, you can turn it off for an entire server.

Maximum disk savings are realized when the following options are selected during INSTALL:

- Compression is enabled.
- Suballocation is enabled.
- Disk allocation blocks are large (16 KB or larger).

File Compression Options

Although file compression is set at the volume level, you control file compression with SET commands or the SERVMAN utility; the settings are stored in the AUTOEXEC.NCF file.

The following table shows the SET commands for managing file compression. These settings apply to all files and directories on the volume.

[illegible]	Explanation
Compression Daily Check Stop Hour	This option specifies when the system should stop looking for unopened files. The default is 6 (6:00 am). Automatic searching for unopened files stops until the starting hour on the next day.
Compression Daily Check Starting Hour	This option specifies when the server should start searching for unopened files. Hours are specified in military time. Default is 0 (midnight).
Minimum Compression Percentage Gain	A file must be at least this specified percentage smaller after compression or the system leaves it uncompressed.
Enable File Compression	If you enter SET Enable File Compression=OFF, file compression is still enabled for the volume, but no files will be compressed. The setting applies to the entire server. Default is ON.
Maximum Concurrent Compressions	This option specifies the number of volumes that can be compressing files on a server at the same time. Default is 2. Novell recommends not changing this setting because increasing the number slows compression down considerably.
Convert Compressed to Uncompressed Option=1	This option specifies how the server stores a compressed file after uncompressing it. Option 0 always leaves the file compressed, 1 leaves it compressed after a single access (within the "untouched" period), and 2 always leaves it uncompressed.
Uncompress Percent Disk Space Free To Allow Commit=10(%)	This option specifies the percentage of free disk space required on a volume before an uncompressed file can be committed to disk. This prevents newly uncompressed files from filling up the volume.
Uncompress Free Space Warning Interval=31 minutes 18.5 seconds	This number specifies the interval for displaying warning alerts when the volume has insufficient free disk space for uncompressed files. To turn off the alerts, set it to 0.
Deleted File Compression Option	This option specifies how the server handles unpurged deleted files: 0=DON'T, 1=Compress next day, 2=Compress immediately.
Days Untouched Before Compression	This option specifies when an unmodified file should be compressed. The system automatically compresses files that have not been accessed for the number of days specified. Default is 7.

Table 9-3: SET Parameters for File Compression

Setting File Compression Attributes

You can set attributes for specific files and directories with the FLAG command (you can also use FILER or the NetWare Administrator). The attributes are Immediate Compress (IC) and Don't Compress (DC).

The Immediate Compress (IC) flag starts compression of the selected files immediately, regardless of the time set for daily compression. If the SET parameter ENABLE FILE COMPRESSION has been set to OFF, compression on these files will be queued until the parameter is reset to ON. You may notice an extremely high processor utilization rate while compression occurs.

The Don't Compress (DC) flag prevents compression of a particular file or of all the files in a subdirectory.

If you group files you want compressed on one volume and files you do not want compressed on another volume, you will not have to control compression at this level.

For example, if you want the files in the subdirectory LETTERS to be compressed each time a file is opened, type the following:

```
FLAG SYS:DOC\LETTERS IC <Enter>
```

If you have an applications directory that contains files or subdirectories that you do not want compressed, type the following:

```
FLAG SYS:APPS\ DC /S <Enter>
```

Viewing Compression Statistics

You can view file compression statistics using the NDIR, FILER, and NetWare Administrator utilities.

NDIR

Use NDIR to view the following statistics:

- Compression statistics for an entire volume: **NDIR [*vol*:] /VOL**
- Compression statistics for a file or directory: **NDIR [*path*] /COMP**

FILER

Select the following options in FILER to view statistics on the volume:

1. Select **View Volume Information.**
2. Select **Statistics.**

To view compression statistics on individual files in FILER, complete the following steps:

1. Select **Manage Files and Directories.**
2. Select an individual file.
3. Select **View/Set file information.**

NetWare Administrator

Complete the following steps to view statistics with NetWare Administrator:

1. Select the Volume object.
2. Select **Object...Details**.
3. Select the **Statistics** page button.

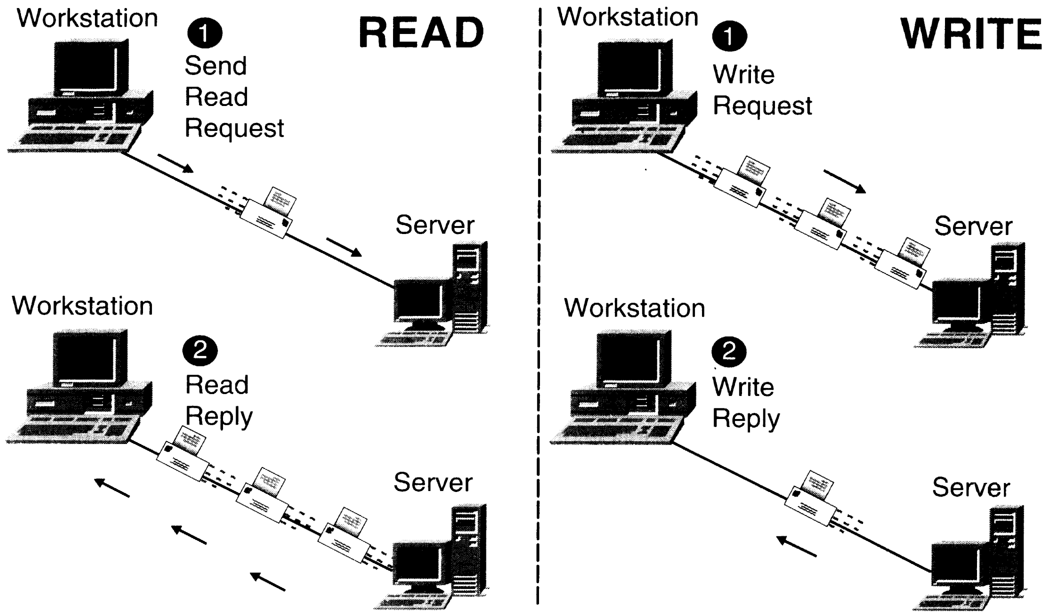

Figure 9-16

Packet Burst Protocol

The Packet Burst™ protocol is a new NetWare Core Protocol™ (NCP) technology that expedites large file reads and writes. The Packet Burst protocol increases the efficiency of client/server communications by allowing workstations to submit a single file read or write request and receive up to 64 KB of data without submitting another request.

Over wide area network (WAN) links, standard NCP requests and responses can result in performance degradation due to the "ping pong" effect. As each station waits for a corresponding acknowledgment, the line is tied up and no other traffic is transmitted.

The delay is further compounded when large files are transmitted one packet at a time. Since the maximum packet size for NetWare routers is 512 bytes, even small files require exchange of numerous packets to fulfill a request.

The Packet Burst protocol can boost performance 10 to 300 percent, depending on factors such as memory, bandwidth, and file sizes.

In NetWare 4, Packet Burst is automatically enabled at both the workstation and the server.

How the Packet Burst Protocol Works

The Packet Burst protocol allows a client to issue a single read or write request for blocks of data up to 64 KB in size. The data is partitioned into packets and these packets are transmitted back-to-back without requiring individual acknowledgment.

Burst Gap Time and Window Size

The Packet Burst protocol uses a *burst gap time* and a *burst window size* to determine how much data is exchanged per *burst* transaction. The *burst gap time* is a time delay requested by the clients before individual packets are placed back-to-back on the media. This *gap time* is used to prevent fast servers from overrunning the client's buffers. The Packet Burst *window size* refers to the number of frames or packets contained in a single burst.

With the Packet Burst protocol, the number of packets in the window is variable, up to a theoretical maximum value of 128 (64 KB divided by 512-byte packets). The window size is also variable by workstation, depending on how many buffers the workstation has available, and by the constant monitoring of the line capacity by the NetWare DOS Requester™. The NetWare DOS Requester renegotiates the parameters without user intervention.

Dropped Packets Recovery

When network traffic is heavy enough that packets are dropped in transmission, the NetWare DOS Requester decreases the window size to minimize packet loss. Any packets that are lost are retransmitted individually; the entire *burst* does not need to be retransmitted.

The Packet Burst protocol monitors dropped packets and retransmits only the missing packets. When Packet Burst-enabled servers or clients transfer data to servers or clients that do not have Packet Burst enabled, the data defaults to normal NCP (one-request/one-response) mode.

The Packet Burst Connection

A workstation sets up a Packet Burst connection with a NetWare server at attach/login time. Once established, the connection remains for the duration of the attachment. If the NetWare DOS Requester fails to make a Packet Burst connection during attach/login, it uses normal NCPs to do the work. By default, Packet Burst is enabled for all workstations.

Packet Burst is established individually for each connection to different file servers. It is therefore possible for a workstation running the NetWare DOS Requester to be attached to a NetWare 4 server using Packet Burst and to a NetWare 3.11 server using a standard NCP connection.

Upon loading, the NetWare DOS Requester first determines whether the workstation has enough memory for the requested number of Packet Burst buffers. If sufficient memory is available, the NetWare DOS Requester initiates a Packet Burst connection to the server.

During this connection process, the NetWare DOS Requester negotiates maximum burst sizes, just as it negotiates packet size, with each connected server.

Once a Packet Burst connection is established between a workstation and a server, the NetWare workstation automatically uses the Packet Burst service whenever an application makes a read or write involving more than 512 bytes of data. Applications do not have to be Packet Burst aware.

Modifying Packet Burst Protocol Performance

In previous versions of NetWare, Packet Burst had to be enabled both at the workstation and at the server. In NetWare 4, Packet Burst is automatically enabled. It cannot be disabled at the server, but it can be disabled on a client-by-client basis.

To enable Packet Burst, add the following line to the NetWare DOS Requester section of the workstation NET.CFG file:

```
PB BUFFERS = n
```

The variable *n* can be any number from 0 to 10. Setting *n* = 0 disables Packet Burst. The safest approach is to set the number of buffers low (2, for example). Increasing the number of Packet Burst buffers does not always result in better performance. Performance can even be degraded if the hardware is unable to handle the additional memory requirement.

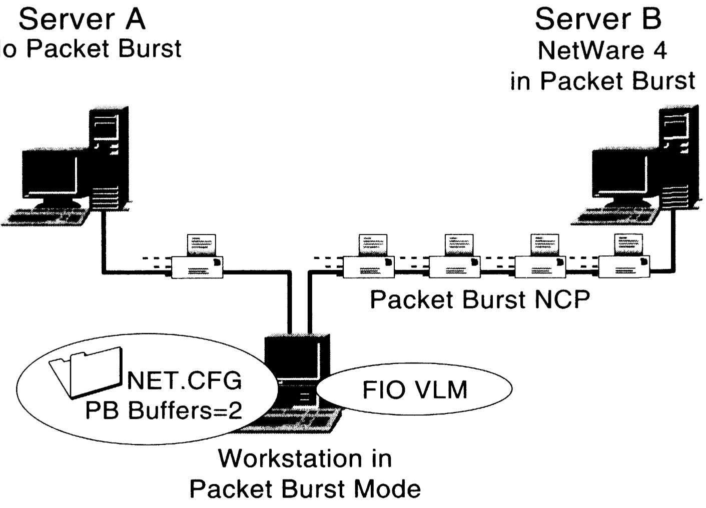

Figure 9-17: Enabling Packet Burst

Large Internet Packets (LIPs)

Some network architectures, such as Ethernet and token ring, allow larger packets to be sent over the network. The Large Internet Packet (LIP) feature enhances throughput on networks with routers by allowing you to specify a large packet size.

On networks without LIP, the workstation and server negotiate to determine the maximum packet size they will exchange. The maximum packet size depends on the packet receive buffer sizes set on both the workstation and the server. You can increase the size on both, but if the server detects routers on the network, the negotiated packet exchange defaults to a 512-byte data packet size.

Enabling LIP causes the server to ignore the router check during packet size negotiation. This allows data packets larger than 512 bytes to be exchanged, as long as the routers are configured to allow data packets larger than 512 bytes.

Just as with Packet Burst, LIP is automatically enabled at both the workstation and the server.

Enabling Routers for Larger Packets

Before LIP can allow larger packets to be exchanged through routers, you must configure the routers to handle larger packet sizes. To configure NetWare routers for larger packet sizes, use the following SET parameter:

> **SET MAXIMUM PHYSICAL RECEIVE PACKET SIZE** = *packet size value*

You can only use this parameter in the STARTUP.NCF file. The change takes effect when the server is restarted. The default value is 4202.

The maximum packet sizes supported by NetWare are as follows:

Topology	Packet Size
Token Ring 16Mbps	4202
Token Ring 4Mbps	4202
Ethernet	1514
ARCnet	4202

Table 9-4: Maximum Packet Sizes

This Maximum Packet Size parameter must be set in all routers on your internetwork. Setting the routers to handle larger packet sizes allows servers and workstations to negotiate the largest packet sizes possible (the smaller of the workstation or server packet sizes).

ARCnet segments will not handle LIPs because of their inherent size limitation.

LIP and Packet Burst Protocol

When LIP is used in conjunction with the Packet Burst protocol, several large packets can be sent consecutively without having to wait for the standard request/response dialogue of normal packet exchanges. This increases data transmission speed while reducing the number of packets exchanged.

Written Exercise 9-2: Interpreting MONITOR Information

Scenario

The following screens show a variety of typical situations you might encounter on your network. Assess each screen to evaluate the potential impact on system performance and consider what actions you might take.

Procedure

Each screen is followed by questions relating to its contents. Write your answers in the space provided. Be prepared to discuss them.

The following figure is an example of a MONITOR Statistics screen. Evaluate the statistics; then answer the questions that follow the figure.

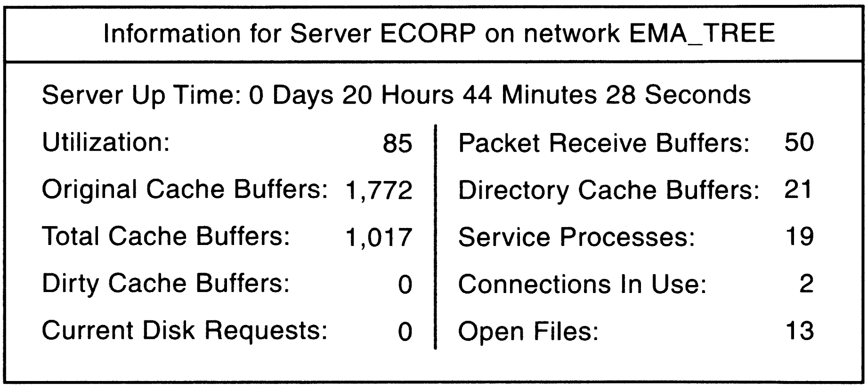

```
            Information for Server ECORP on network EMA_TREE

  Server Up Time: 0 Days 20 Hours 44 Minutes 28 Seconds
  Utilization:                  85 | Packet Receive Buffers:    50
  Original Cache Buffers:    1,772 | Directory Cache Buffers:   21
  Total Cache Buffers:       1,017 | Service Processes:         19
  Dirty Cache Buffers:           0 | Connections In Use:         2
  Current Disk Requests:         0 | Open Files:                13
```

1. What items on the information screen above need attention?

   ut = 85

2. Suggest actions to take for each item.

   interrupts

The following figure is an example of a Cache Utilization screen. Evaluate the statistics; then answer the questions that follow the figure.

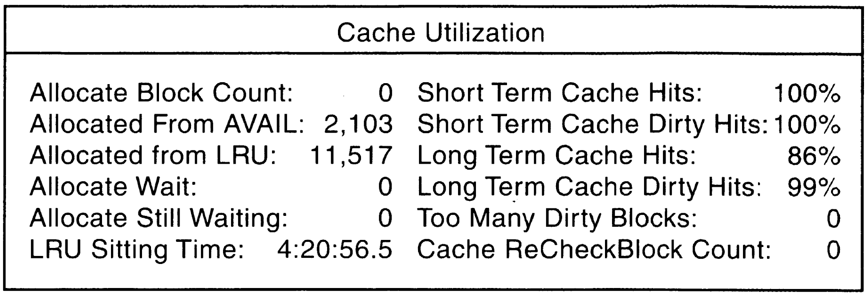

```
                          Cache Utilization

Allocate Block Count:          0  Short Term Cache Hits:         100%
Allocated From AVAIL:      2,103  Short Term Cache Dirty Hits:   100%
Allocated from LRU:       11,517  Long Term Cache Hits:           86%
Allocate Wait:                 0  Long Term Cache Dirty Hits:     99%
Allocate Still Waiting:        0  Too Many Dirty Blocks:            0
LRU Sitting Time:      4:20:56.5  Cache ReCheckBlock Count:         0
```

1. What items on the information screen above need attention?

   L.t.c.h.

2. Suggest actions to take for each item.

   : modules

The following figure is an example of network board statistics. Evaluate the statistics; then answer the questions that follow the figure.

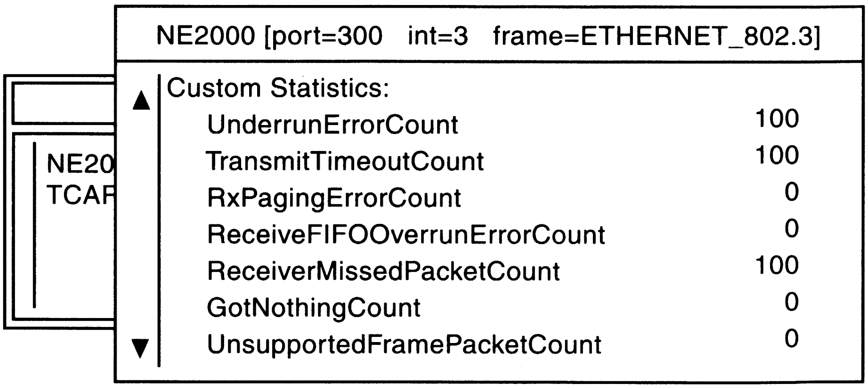

```
NE2000 [port=300   int=3   frame=ETHERNET_802.3]

NE20   Custom Statistics:
TCAF     UnderrunErrorCount                        100
         TransmitTimeoutCount                      100
         RxPagingErrorCount                          0
         ReceiveFIFOOverrunErrorCount                0
         ReceiverMissedPacketCount                 100
         GotNothingCount                             0
         UnsupportedFramePacketCount                 0
```

1. What items on the information screen above need attention?

? x 100

2. Suggest actions to take for each item.

type netwerkkaart OK ?

The following figure gives polling and processor utilization statistics. Evaluate the statistics; then answer the questions that follow the figure.

Name	Time	Count	Load
Polling	1,124,257	18	96.33%
Total Sample Time:	1,179,997		
Histogram Overhead Time:	12,997	(1.10%)	
Adjusted Sample Time:	1,167,000		

1. What does the figure above tell you?

2. Would you initiate any actions based on this information?

3. If polling percentage was low, what value would you expect to see in processor utilization?

4. If processor utilization is consistently over 80 percent, what might you do?

The following figure provides memory utilization statistics. Evaluate the statistics; then answer the questions that follow the figure.

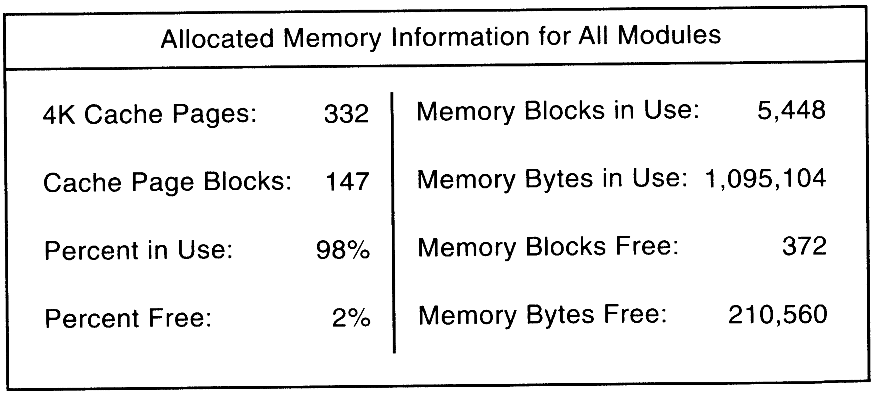

```
              Allocated Memory Information for All Modules

 4K Cache Pages:        332 | Memory Blocks in Use:        5,448
                            |
 Cache Page Blocks:     147 | Memory Bytes in Use:     1,095,104
                            |
 Percent in Use:        98% | Memory Blocks Free:            372
                            |
 Percent Free:           2% | Memory Bytes Free:         210,560
```

1. What items on the information screen above need attention?

2. Suggest actions to take for each item.

Written Exercise 9-3: Optimizing Network Performance

Scenario

The Research department is increasingly writing many small files to disk for a large database on your server. Users from other divisions within the organization are accessing this database and are causing additional traffic across the network.

Procedure

The following questions are based on the scenario. Write your answers in the space provided. Be prepared to discuss them.

1. What impact might a high-end device driver/controller combination have on this situation?

2. Would you consider altering the Turbo FAT Re-Use Wait Time parameter in this situation? Why?

3. Would implementing the Packet Burst protocol have a positive effect on network performance for this scenario? Why?

4. Would implementing LIP have a positive effect on network performance for this scenario? Why?

Summary

One of the most challenging and rewarding responsibilities of a network administrator is optimizing performance on the network.

You should look at several areas when you assess the efficiency of your network or encounter a performance problem. These areas include the following:

- The performance statistics provided through MONITOR, including cache utilization, processor utilization, memory utilization, and resource utilization, give you valuable insight into the efficiency of your network.
- Server packet and buffer parameters can be modified to enhance server performance.
- You can change the SAP and RIP parameters to counter bottlenecks caused by large amounts of SAP and RIP broadcast traffic.
- Network board statistics may indicate a performance problem that can be corrected by upgrading the network board.
- Performance is affected by the disk drivers and controllers used by the server. You can gain significant performance improvements with high-end disk driver/controller configurations.
- The turbo FAT indexing feature of NetWare 4 optimizes indexing and accessing very large files. You can modify the turbo FAT parameters to suit your network needs.
- The Packet Burst protocol speeds the transfer of multiple-packet NCP reads and writes and can boost performance 10 to 300 percent.
- Your network architecture may allow you to send larger packets over the network. If so, you can use LIP, which enhances throughput on networks with routers. You can specify a large packet size to gain this benefit.
- You can use LIP with the Packet Burst protocol to increase data transmission speed while reducing the number of packets exchanged.

The system achieves memory protection by separating the memory into domains and providing related processes with access to their own memory allocations. NLMs can be coded for NetWare 4 and tested in a domain established for their execution where they cannot crash the NOS. This domain is created with DOMAIN.NLM.

Notes

APPENDIX A Case Company

Following is an organizational chart for the case company, Enterprise Marketing Association (EMA).

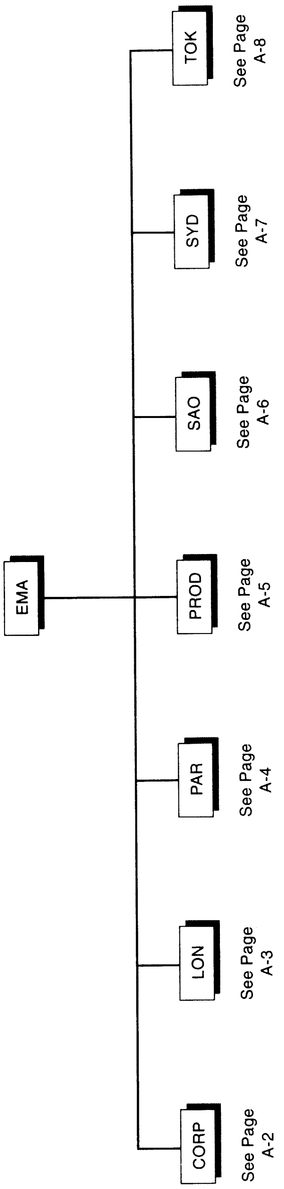

Appendix A

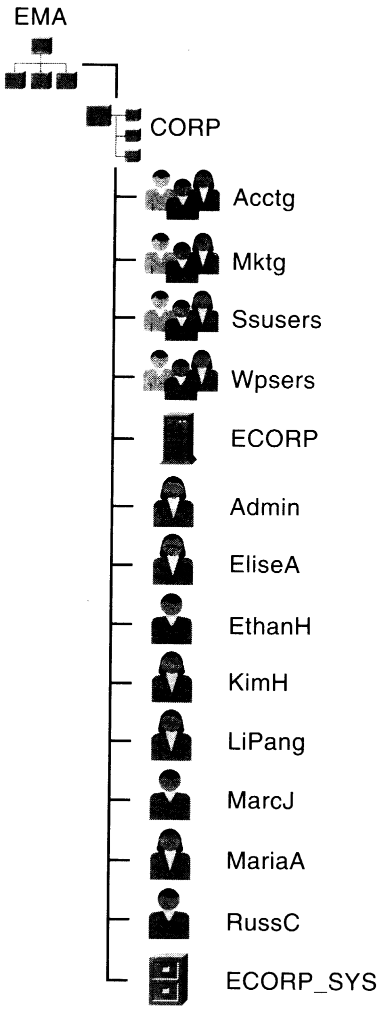
EMA
CORP
Acctg
Mktg
Ssusers
Wpsers
ECORP
Admin
EliseA
EthanH
KimH
LiPang
MarcJ
MariaA
RussC
ECORP_SYS

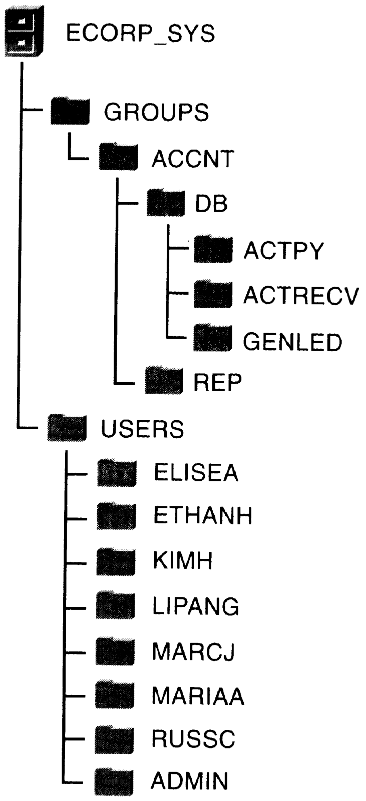
ECORP_SYS
GROUPS
ACCNT
DB
ACTPY
ACTRECV
GENLED
REP
USERS
ELISEA
ETHANH
KIMH
LIPANG
MARCJ
MARIAA
RUSSC
ADMIN

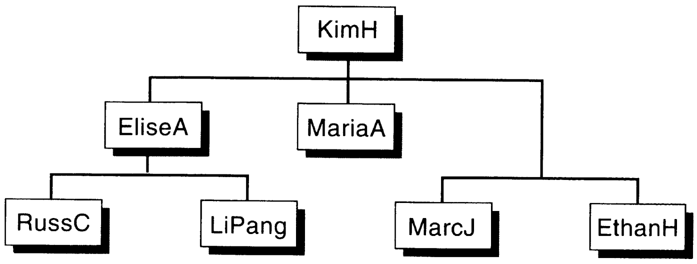
KimH
EliseA
MariaA
RussC
LiPang
MarcJ
EthanH

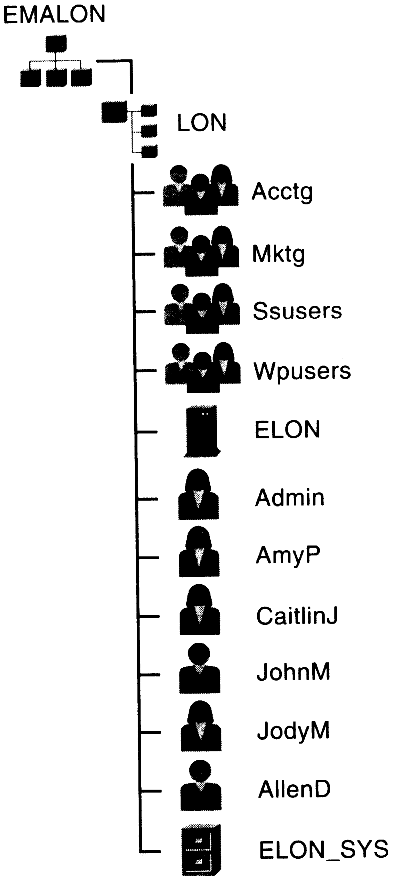
EMALON
LON
Acctg
Mktg
Ssusers
Wpusers
ELON
Admin
AmyP
CaitlinJ
JohnM
JodyM
AllenD
ELON_SYS

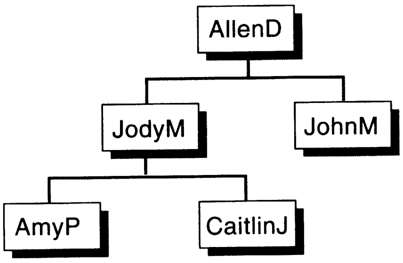
AllenD
JodyM
JohnM
AmyP
CaitlinJ

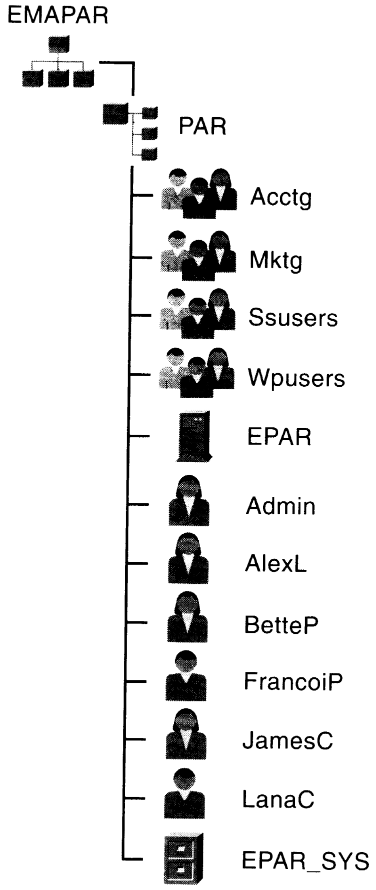
EMAPAR
PAR
Acctg
Mktg
Ssusers
Wpusers
EPAR
Admin
AlexL
BetteP
FrancoiP
JamesC
LanaC
EPAR_SYS

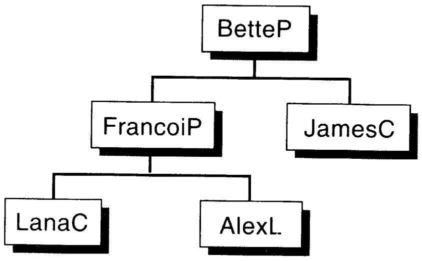
BetteP
FrancoiP
JamesC
LanaC
AlexL

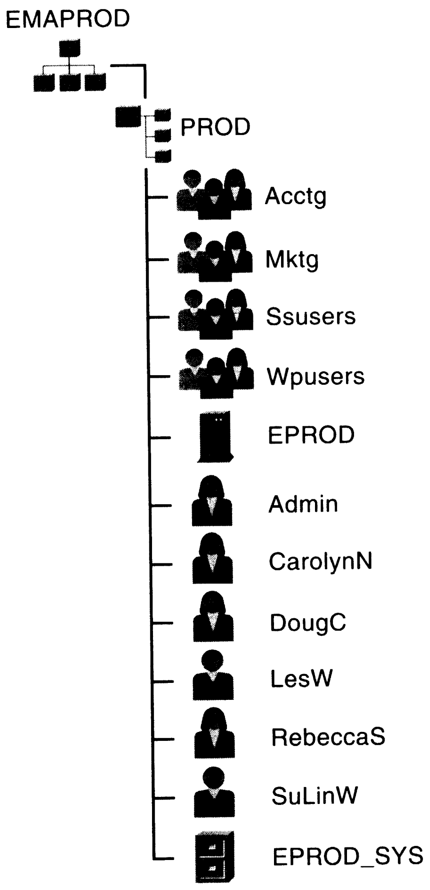
EMAPROD
PROD
Acctg
Mktg
Ssusers
Wpusers
EPROD
Admin
CarolynN
DougC
LesW
RebeccaS
SuLinW
EPROD_SYS

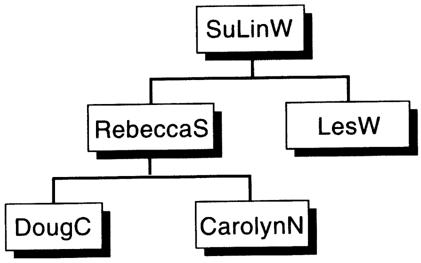
SuLinW
RebeccaS
LesW
DougC
CarolynN

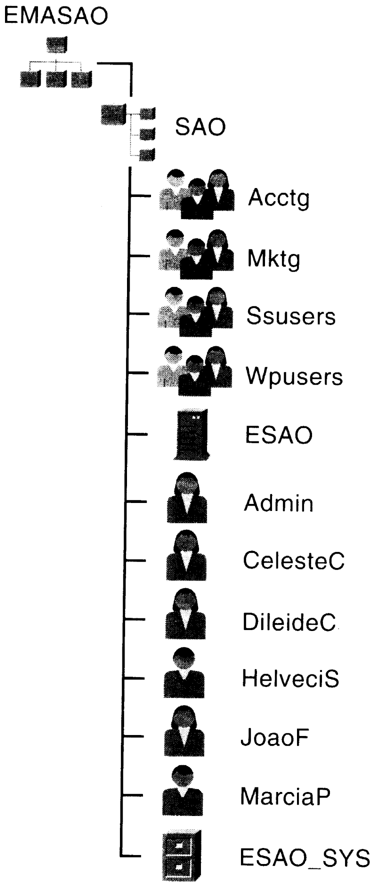
EMASAO
SAO
Acctg
Mktg
Ssusers
Wpusers
ESAO
Admin
CelesteC
DileideC
HelveciS
JoaoF
MarciaP
ESAO_SYS

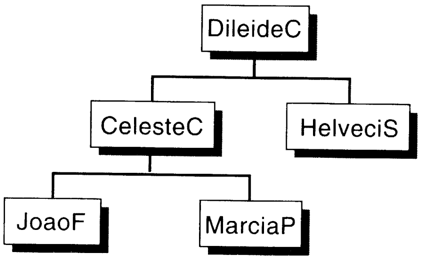
DileideC
CelesteC
HelveciS
JoaoF
MarciaP

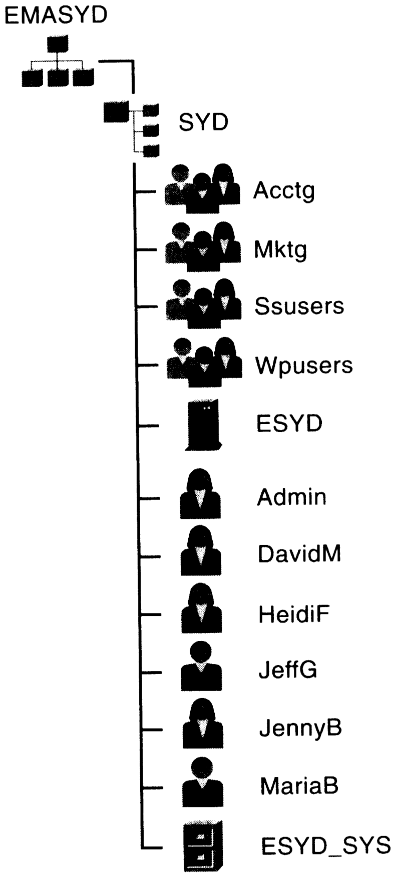
EMASYD
SYD
Acctg
Mktg
Ssusers
Wpusers
ESYD
Admin
DavidM
HeidiF
JeffG
JennyB
MariaB
ESYD_SYS

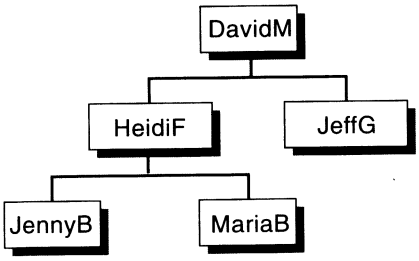
DavidM
HeidiF
JeffG
JennyB
MariaB

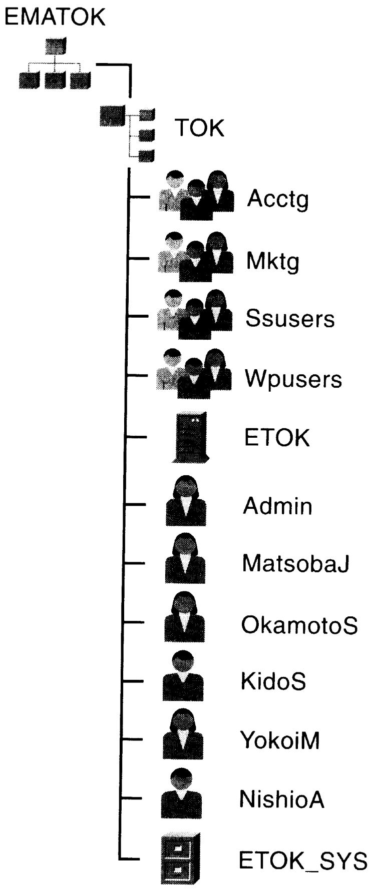
EMATOK
TOK
Acctg
Mktg
Ssusers
Wpusers
ETOK
Admin
MatsobaJ
OkamotoS
KidoS
YokoiM
NishioA
ETOK_SYS

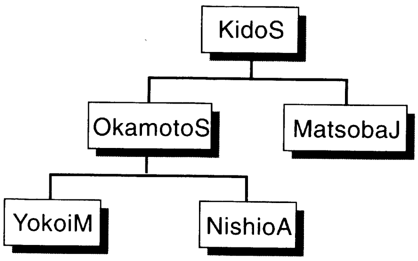
KidoS
OkamotoS
MatsobaJ
YokoiM
NishioA

APPENDIX B Selected Application Notes

This appendix contains the following *Novell Application Notes*:

May 1994	The Functions and Operations of the NetWare DOS Requester 1.1
June 1994	Compression and Suballocation in NetWare 4

These *Application Notes* are included with permission from Novell Research.

Notes

MAY 1994

NOVELL® RESEARCH

The Functions and Operations of the NetWare DOS Requester 1.1

John Froelich
Software Test Engineer Associate
VLM Development Team

Ed Liebing
Senior Editor
Systems Research Department

Brad Young
Product Support Engineer
Worldwide Customer Services and Support

This AppNote explains the DOS Requester's theory of operations as well as some noteworthy differences between the NetWare Shell and the DOS Requester. Next, the AppNote looks at how the DOS Requester initializes and how it performs a read request. Finally, example settings of workstation NET.CFG files help explain performance and memory considerations.

Trademarks

Novell, the N-Design, and NetWare are registered trademarks, and NetWare Directory Services, NDS, NetWare DOS Requester, NetWare Loadable Module, NLM, Virtual Loadable Module, and VLM are trademarks of Novell, Inc.

All other product names mentioned are trademarks of their respective companies or distributors.

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of these Application Notes (AppNotes) or of any of the third-party products discussed in the AppNotes. Novell reserves the right to revise these AppNotes and to make changes in their content at any time, without obligation to notify any person or entity of such revisions or changes. These AppNotes do not constitute an endorsement of the third-party product or products that were tested. Configuration(s) tested or described may or may not be the only available solution. Any test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state, or local requirements. Novell does not warranty products except as stated in applicable Novell product warranties or license agreements.

Novell, Inc.
122 East 1700 South
Provo, Utah 84606 USA

Contents

Acknowledgements

We would like to thank Bart Reese, Ian Stiles, and Jay Sevison for their assistance with the technical aspects of this AppNote.

Introduction

In April of 1993, Novell Research published its special NetWare 4.0 edition of *NetWare Application Notes.* The issue contained a rather lengthy write-up on the new DOS Requester 1.0 that shipped with NetWare 4. Since then, the DOS Requester has gone through a number of enhancements to improve its overall performance and functionality, and has now been upgraded to version 1.1.

This AppNote takes the next step in explaining how the DOS Requester works. The Introduction first lists the new features that come with the DOS Requester 1.1. The AppNote then discusses the DOS Requester's theory of operations and how the NetWare Shell and the DOS Requester work with DOS. We then examine how the DOS Requester initializes and how it performs a read request.

Finally, we look at three NET.CFG files to show some different ways you can configure the DOS Requester. The sample files emphasize performance optimization, memory optimization, and a balance between the two.

What's New in the DOS Requester 1.1

The following is a list of the new features added to the DOS Requester since the release of NetWare 4.01:

- Connection timeout optimization
- Personal NetWare client functionality
- Network Management Responder (NMR.VLM)
- Optimized Packet Burst over WAN and other slow links
- IPX short-lived socket support
- New NET.CFG parameters for increased flexibility
- Performance optimization for WAN connections
- Overall optimized performance and software compatibility

The DOS Requester Theory of Operations

A VLM is a modular program that contains logically grouped features in order to perform logically grouped functions. For example, transport-related functions such as sending a packet and receiving a packet fit logically into the IPXNCP.VLM module.

The DOS Requester architecture includes various categories of services and components. Individual architectural pieces include the following:

- The VLM Manager (VLM.EXE)
- Multiplexor VLMs
- Child VLMs
- "Standard" VLMs

VLM.EXE is a memory manager that coordinates and controls the VLMs at each of the Requester layers. VLM.EXE loads and unloads the VLM modules and handles memory services for the modules.

A *multiplexor* VLM module provides a common API interface to dissimilar child VLM modules. A *child* VLM module contains the logical grouping of functionality and loads prior to its multiplexor module. A *standard* VLM module offers functionality that does not fall into the category of a multiplexor or a child VLM module. An example of a standard module is AUTO.VLM, which offers automated reconnection capabilities.

The DOS Requester interacts directly with the ODI network communications modules and with DOS. All of these pieces comprise the overall networking environment. This concept is graphically represented in both task-switching and non-switching environments in Figure 1.

Figure 1: The DOS Requester interacts with DOS in both task-switching and non-switching environments.

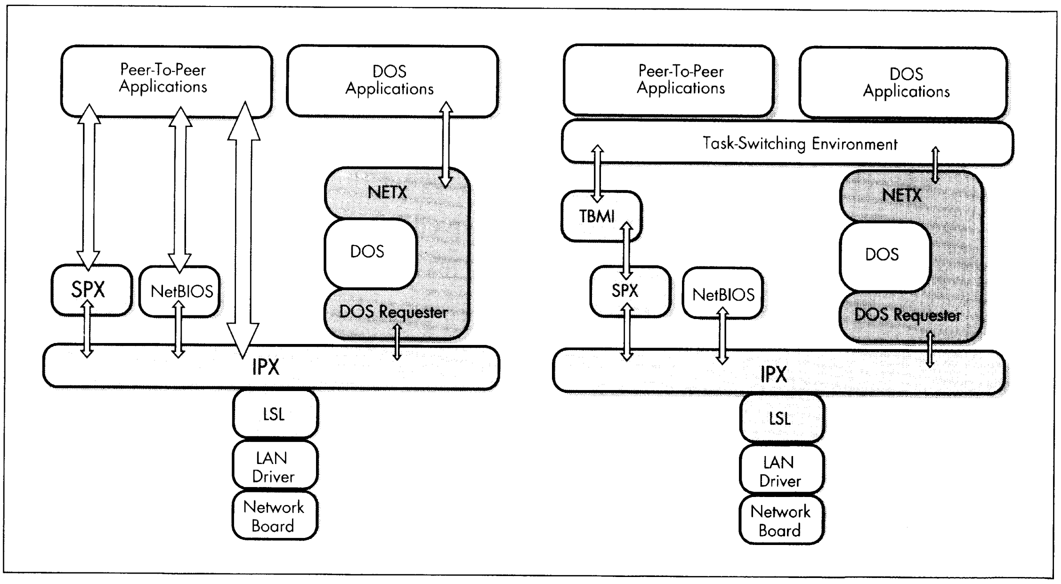

The VLM Manager

The VLM Manager oversees the operations of the VLM modules and handles memory management services for them. Applications making calls to the DOS Requester have *all* their calls routed through VLM.EXE, which directs the requests to their proper destinations. VLM.EXE also ensures that replies return to their respective callers.

The VLM Manager ensures that the API calls between other modules are properly routed. The VLM Manager must therefore know if the modules required to satisfy an API call are loaded into memory. It must police VLM modules that call other modules and handle asynchronous calls to the VLMs and their related functions. Even when a child module calls its multiplexor, the call goes to the VLM Manager, which then calls the specified multiplexor.

Figure 2: The VLM Manager ensures that API calls between other modules are properly routed.

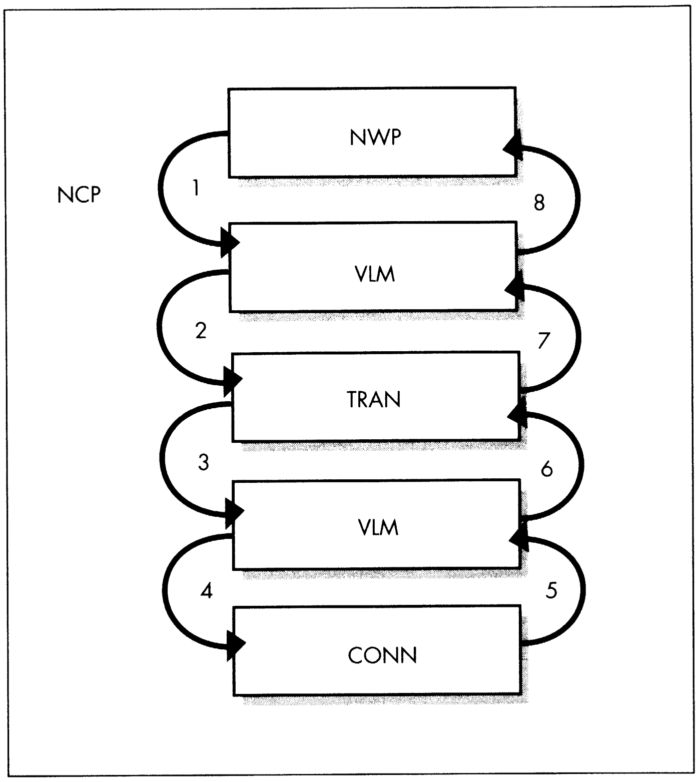

The VLM Manager provides the API interface to call a function and a VLM module by its assigned number. For example, when a request calls for a module, the caller passes the caller's ID (0 for an application, a non-zero for a VLM), the destination VLM's ID, and the function number.

Modules that are not part of the DOS Requester model can also use the VLM Manager as a TSR memory manager. NMR.VLM, the NetWare Management Responder module, is a good example of a module that uses VLM.EXE in this way.

The VLM Manager decides whether the VLM modules use extended memory, expanded memory, conventional memory, or any memory type supported. This frees the VLM modules from memory model decisions and concerns.

The VLM Manager handles *transient block swapping* and uses it to swap the transient segments of the VLM modules into and out of XMS or EMS memory. Each VLM has a portion of its code (called the global segment) that must be maintained in conventional memory, such as interrupt handlers, ESRs (Event Service Routines), buffers passed as pointers, and asynchronous event handlers. Lesser-used function calls are located in the transient segment and are placed into XMS or EMS memory (if you are loading one of these memory managers). These are swapped back into the transient swap block in conventional memory or in upper memory when the module's functions need to be executed. The module is then swapped back out to XMS or EMS memory when another module is placed in the transient swap block. The effect of swapping the transient portion back out is that transient data that may have changed is preserved.

Figure 3: The VLM Manager uses module swapping to swap transient portions of VLM modules into and out of memory.

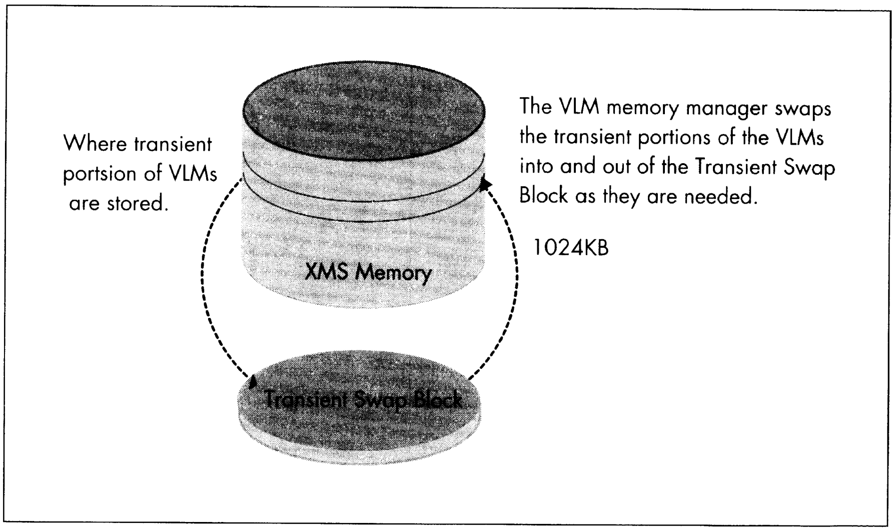
Where transient portsion of VLMs are stored.
The VLM memory manager swaps the transient portions of the VLMs into and out of the Transient Swap Block as they are needed.
1024KB
XMS Memory
Transient Swap Block

Transient block swapping allows for a much smaller conventional memory footprint, with much of the lesser-used code residing in a swapped-out state (usually XMS or EMS) until it is required. Swapping, therefore, helps balance UMB (upper memory block) and conventional memory use with performance.

The VLM Manager is also responsible for load time configuration APIs. For instance, a user may want CONN.VLM to have 16 connections in its connection table, or FIO to support a larger cache size and a larger number of buffers. In these instances, the VLM Manager reads pertinent information from the NET.CFG file, returns the configuration data to the specific VLM, and expands or contracts the VLM accordingly.

Multiplexors

Multiplexors can be understood as "parent" VLMs and route calls to their registered child VLMs. Multiplexors insulate an application from the idiosyncrasies of a child VLM's services (such as login or attachment variations between Directory Services, bindery-based servers, and Personal NetWare).

The NWP.VLM module acts as the parent multiplexor to NDS.VLM, BIND.VLM and PNW.VLM. If a workstation issues a call to attach to a bindery server, NWP.VLM routes the call to BIND.VLM for processing. On the other hand, if a workstation makes a call to attach to *any* type of server (referred to as a "wildcard" call), NWP.VLM traverses the child VLMs in the order they were loaded until one of the modules can satisfy the call.

Child VLMs

Child VLMs are different implementations of a logical grouping of functionality. Each server type has its own child VLM:

- BIND.VLM for bindery-based/pre-NetWare 4 servers
- NDS.VLM for NetWare 4 Directory Services-based servers
- PNW.VLM for Personal NetWare-based servers

Different implementations of transport protocols can also have individual VLMs. For example, IPXNCP.VLM handles IPX services, and future releases may include a VLM to handle TCP/IP services, with both modules using the TRAN.VLM (the transport protocol) multiplexor.

Note: Although NETX.VLM is the VLM's functional replacement for NETX.COM or NETX.EXE, IPXNCP.VLM is not a replacement for IPXODI.COM. In fact, IPXNCP.VLM requires that IPXODI.COM or IPX.COM be loaded.

The load order of child VLMs can also determine the behavior of "wildcard" calls. If a workstation loads the BIND.VLM module before NDS.VLM, the preferred attachment to a server for the "wildcard" AttachToServer call will be a bindery type as opposed to NDS, even if you are only logging in to a NetWare 4 server.

You also see this type of behavior when using `PREFERRED SERVER` or `PREFERRED TREE` statements in the NET.CFG or when using `/PS=` or `/PT=` parameters when you load VLM.EXE at the command line. If BIND.VLM is loaded first, a `PREFERRED SERVER` parameter overrides a `PREFERRED TREE` parameter. If NDS.VLM is loaded before BIND.VLM, the opposite occurs.

For this reason, the DOS Requester 1.1 contains the NET.CFG parameter `Netware Protocol = <vlm>[, <vlm>...]`, which provides an easy way to maintain these scenarios. For example, `Netware Protocol = BIND,NDS` places bindery services before Directory Services. (You can also exclude any of the child VLM modules of NWP.VLM by not including that module in the list.)

Categories of Services

Each VLM operates at a different protocol layer with a specific category of service. These categories of services form three distinct layers in the Requester architecture:

- The DOS Redirection Layer
- Service Protocol Layer
- Transport Protocol Layer

Modules at a given layer do not need to be concerned with a particular configuration that resides below them.

Figure 4: VLM services form three distinct layers in the Requester architecture.

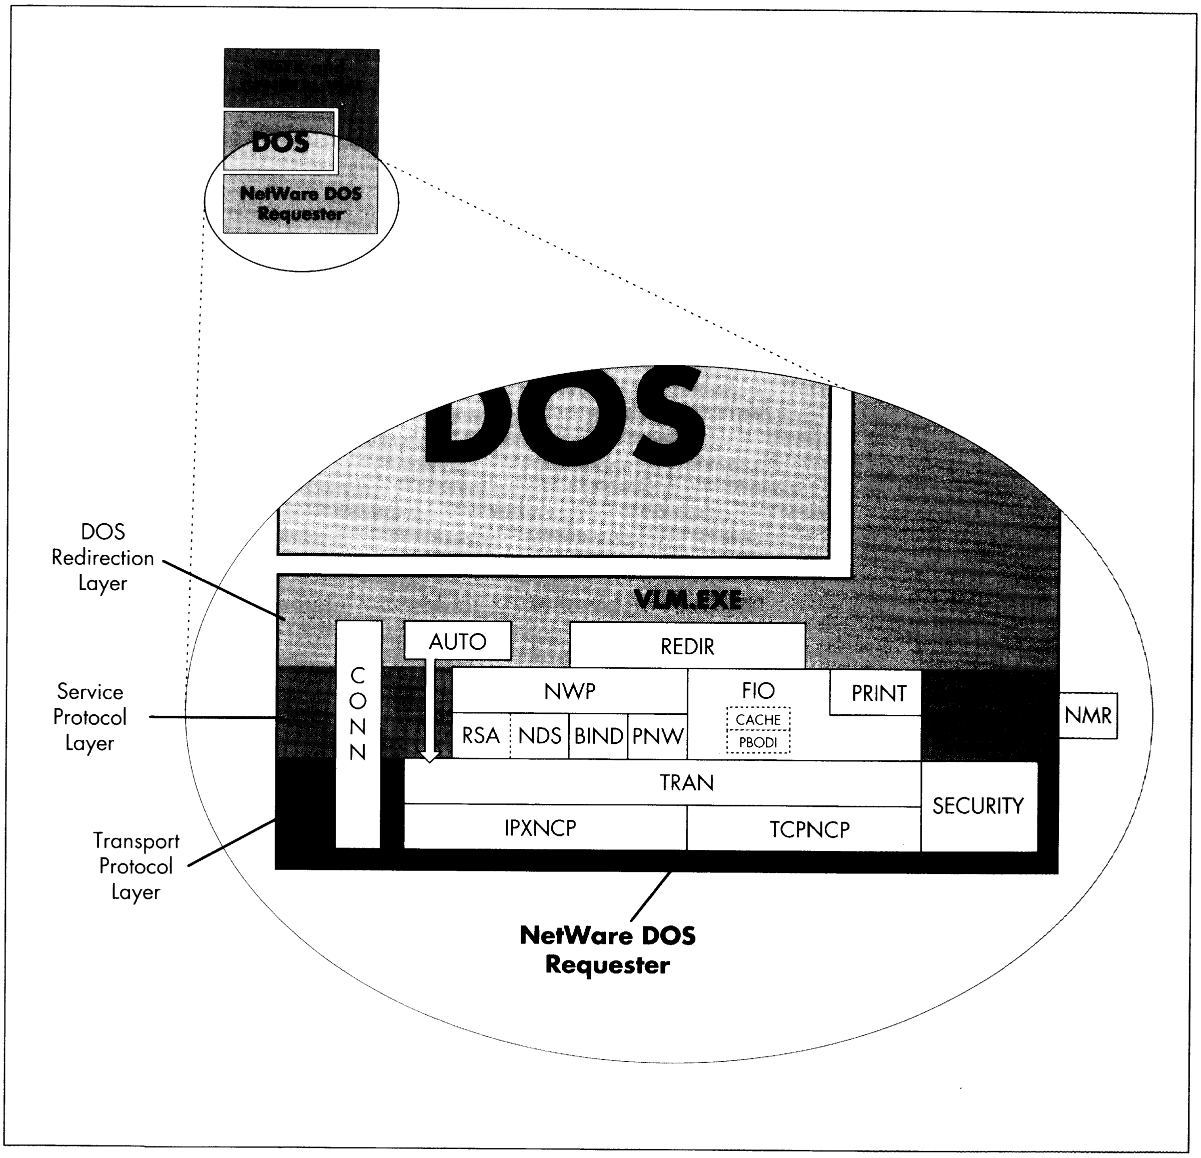

The DOS Redirection Layer. This layer provides DOS redirection services for the DOS Requester with most of the functionality bundled into the REDIR.VLM module.

With the DOS Requester, DOS uses the same mechanism to recognize a network drive as it does to recognize CD-ROM drives. CD-ROM drivers use the DOS Redirector interface to make the drive available to the workstation. Network redirectors use the same interface to make file servers available for clients. REDIR.VLM provides the redirector support for the DOS Requester and makes the NetWare server's resources appear as DOS resources to the user.

Figure 5: The REDIR.VLM module provides redirector support for the DOS Requester.

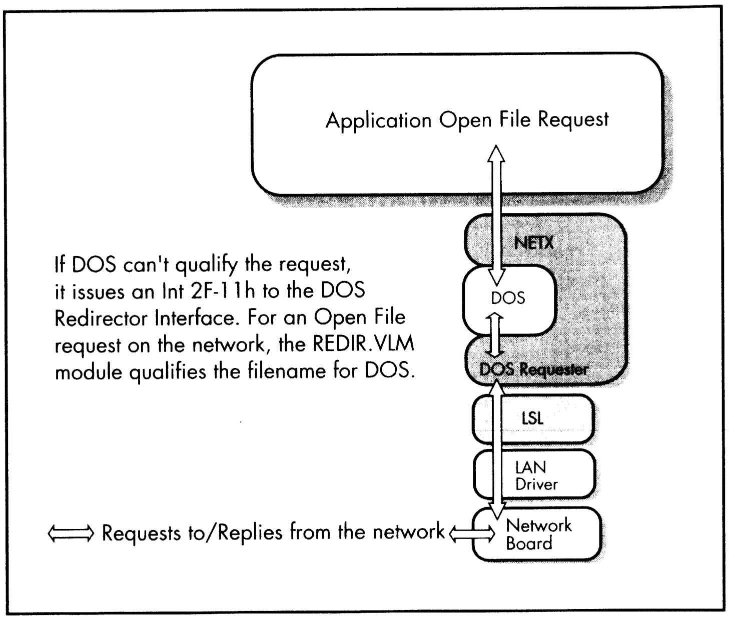
Application Open File Request
NETX
DOS
DOS Requester
LSL
LAN Driver
Network Board
⟺ Requests to/Replies from the network

If DOS can't qualify the request, it issues an Int 2F-11h to the DOS Redirector Interface. For an Open File request on the network, the REDIR.VLM module qualifies the filename for DOS.

The Service Protocol Layer. The Service Protocol Layer includes parallel protocols, file services, and print services. The NWP.VLM module is the NetWare protocol portion for the DOS Requester and consists of the NDS-, bindery-, and Personal NetWare-based server types. The FIO.VLM module is the file transfer portion and handles NetWare file I/O traffic. The print handler portion, the PRINT.VLM module, provides printer redirection services. RSA.VLM is not loaded by default, and is only needed for auto-reconnect services to a directory services-based server. Utilities such as LOGIN.EXE take care of initial authentication services to directory services-based servers.

NWP.VLM is the NetWare protocol multiplexor that handles the network-server implementations, which include the Directory

Service module (NDS.VLM), bindery module (BIND.VLM), or Personal NetWare module (PNW.VLM). Services provided by the NWP.VLM module include establishing and destroying connections, as well as logins and logouts.

The file input/output module (FIO.VLM) also resides within the Service Protocol Layer and provides a basic file transfer protocol. The module's additional capabilities include cached or non-cached read/writes, Packet Burst mode-based read/writes, Large Internet Packet support, and others. The ability to switch these features on or off allows users to choose performance over memory usage, or vice versa. The PRINT.VLM module provides network printing services. Because PRINT.VLM uses FIO.VLM for its file writes, the printing services can be handled in the following ways:

- Non-cached
- Cached
- Redirected via Packet Burst protocol to the server

The FIO.VLM modules gives the PRINT.VLM module added performance by allowing print jobs to use larger cached writes for faster printing.

The Transport Protocol Layer. The Transport Protocol Layer is responsible for maintaining connections, providing packet transmissions and receptions between connections, and providing other transport-specific services. The term *transport-specific* refers to similar services implemented by different transport protocols. The IPXNCP.VLM module uses IPX as its transport protocol. The TRAN.VLM module shields the layers above it from the information handled by the transport protocol modules.

The NetWare DOS Requester architecture allows VLMs of a given layer to coexist and to perform their respective functions without conflict. For example, the TRAN.VLM, IPXNCP.VLM, and SECURITY.VLM modules operate at the Transport Protocol Layer. Similarly, service-protocol VLMs like FIO.VLM and PRINT.VLM operate at the Service Protocol Layer.

The *Connection Table Manager* (CONN.VLM) spans all three Requester functional layers and provides the connection services required by all levels of the Requester model.

How the Shell and the DOS Requester Work with DOS

The old NetWare Shells, NETX.COM and NETX.EXE, XMSNETX.EXE, and EMSNETX.EXE, all run "in front" of DOS. That is, the Shell intercepts interrupts (requests for service) that are directed to DOS and acts on those requests associated with a NetWare resource. If the request is not for the network, the Shell passes the interrupt (request) to DOS.

The NetWare DOS Requester operates differently than the NetWare Shell because the Requester runs "behind" DOS. DOS gives network requests to NetWare through the DOS Redirector Interface (Int 2Fh function 11). This interface is the mechanism (by design) which allows DOS to recognize foreign file systems, such as CD-ROMs and networks, and can in theory be used by any number of redirectors. The NetWare DOS Requester interacts with this interface through the REDIR.VLM module.

DOS calls redirectors when DOS itself cannot service a request. For example, when an application makes a request to DOS, DOS first attempts to qualify the request in order to determine ownership of the requested resource, such as a drive letter, a file handle, or a print device. If DOS determines that it does not own the resource requested, DOS polls the redirectors to allow them to determine ownership among themselves. If multiple redirectors are loaded, each redirector in turn attempts to qualify the request. If a redirector claims ownership to the requested resource, DOS then makes the appropriate calls to that redirector so it can complete the request; otherwise, the application receives an error for the request.

For example, if the workstation is running the DOS Requester and an application makes a DOS Int 21-3Fh "Open File" request, DOS attempts to qualify the request. If DOS can't, it issues an Int 2F-1123h call to the REDIR.VLM module in order to qualify the filename. When DOS receives a successful return code from REDIR, DOS issues an Int 2F-1116h call to open the file. The REDIR.VLM then requests other services within the DOS Requester that forward the call to the file server. The file server returns the resulting information and status codes to the DOS Requester, which then returns the information to the REDIR.VLM module, then to DOS, and then to the application.

If the workstation also loads the NetWare Shell emulator, NETX.VLM, the emulator acts like the NetWare Shell and intercepts Int 21h requests. Requests that are within the range of the old NetWare APIs (B0h through E3h) are routed immediately to the network and bypass DOS entirely. Otherwise, the request goes to DOS for servicing.

NetWare Shell emulation makes it possible for the DOS Requester to run applications that use the old Shell calls. However, the functionality of the NETX.VLM module differs from the NetWare Shell because the Shell intercepts *all* Int 21h calls and if the call involves a NetWare resource, the Shell processes them without any DOS intervention.

Note: Because of the different methods they use to manage resources, the NetWare Shell (NETX.EXE) and the DOS Requester *cannot* coexist and cannot be loaded together.

In order for the NetWare Shell to completely circumvent DOS, it uses its own set of internal resource tables for network file and print services. (For an explanation of how the NetWare Shell and the DOS Requester handle drive mappings, see the "Search Drive Mappings and the PATH Variable" and the "Still Searching . . . " NetNotes in the October 1993 and December 1993 issues of *NetWare Application Notes,* respectively.)

In contrast, the DOS Requester is tightly integrated with DOS and, as a redirector, shares resource tables with DOS. This eliminates the need to maintain redundant tables, decreasing memory requirements for the DOS Requester. However, if you load the NETX.VLM module, some of the NetWare Shell's internal tables are also maintained for backward compatibility.

How the VLMs Initialize

Because the DOS Requester is designed in a modular fashion, it can be set up in a variety of memory configurations. The VLMs can be loaded into extended memory, expanded memory, or conventional memory, and they will use upper memory blocks (rather than conventional memory) when such memory is available. The following is a description of how the DOS Requester accomplishes this.

Figure 6: VLM.EXE relocates its startup code.

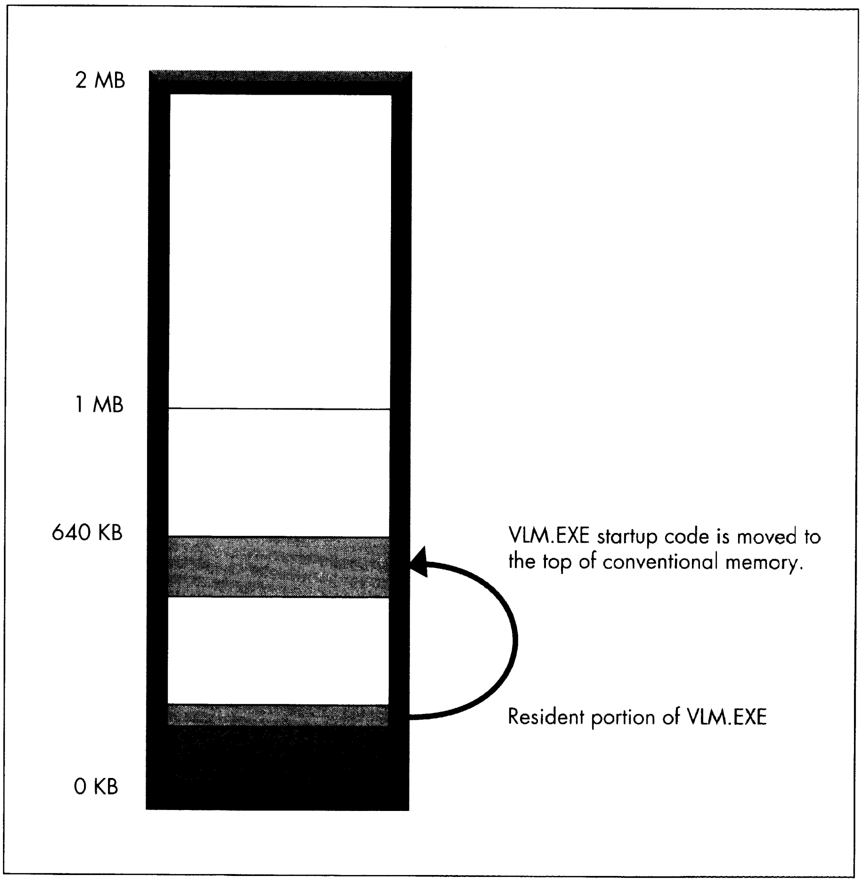
2 MB
1 MB
640 KB
0 KB
VLM.EXE startup code is moved to the top of conventional memory.
Resident portion of VLM.EXE

As shown in Figure 6, during the pre-initialization process, VLM.EXE relocates its startup code at the top of conventional memory so it can load the VLM modules and then remove itself without occupying memory blocks unnecessarily.

At this point, VLM.EXE performs several checks and initializations before loading the startup information. These checks include detecting whether the DOS Requester has already been loaded, whether it has been loaded under a task switcher, and whether the command line or the NET.CFG file contains specific parameter settings.

Figure 7: The VLM Manager checks memory requirements for each module.

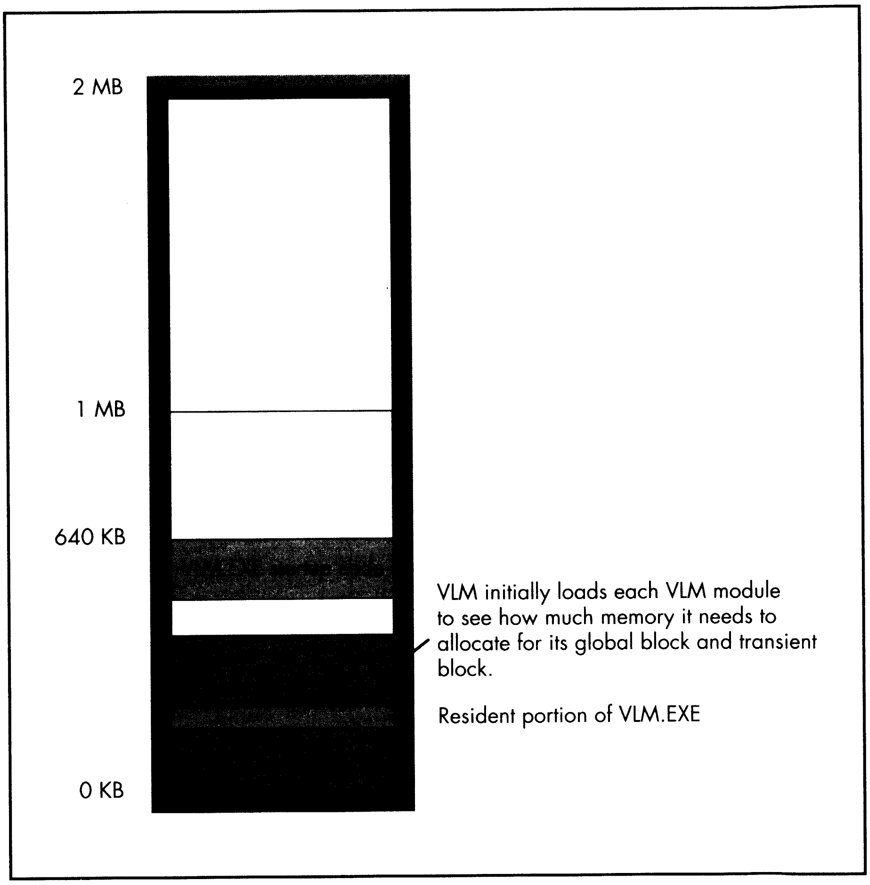
2 MB
1 MB
640 KB
VLM.EXE startup code
0 KB
VLM initially loads each VLM module to see how much memory it needs to allocate for its global block and transient block.
Resident portion of VLM.EXE

The VLM manager is particularly interested in which VLMs are loading. VLMs are load-order dependent and follow the hard-coded defaults built into VLM.EXE unless you specify otherwise in the NET.CFG.

VLM.EXE performs a pre-initialization process by temporarily loading each VLM module and passing control to the module's initialization code. The module reports its ID and global, transient, and startup memory segment sizes. As this step happens, you see a series of dots displayed on the screen.

Figure 8: The VLM Manager takes advantage of XMS or EMS if either is present.

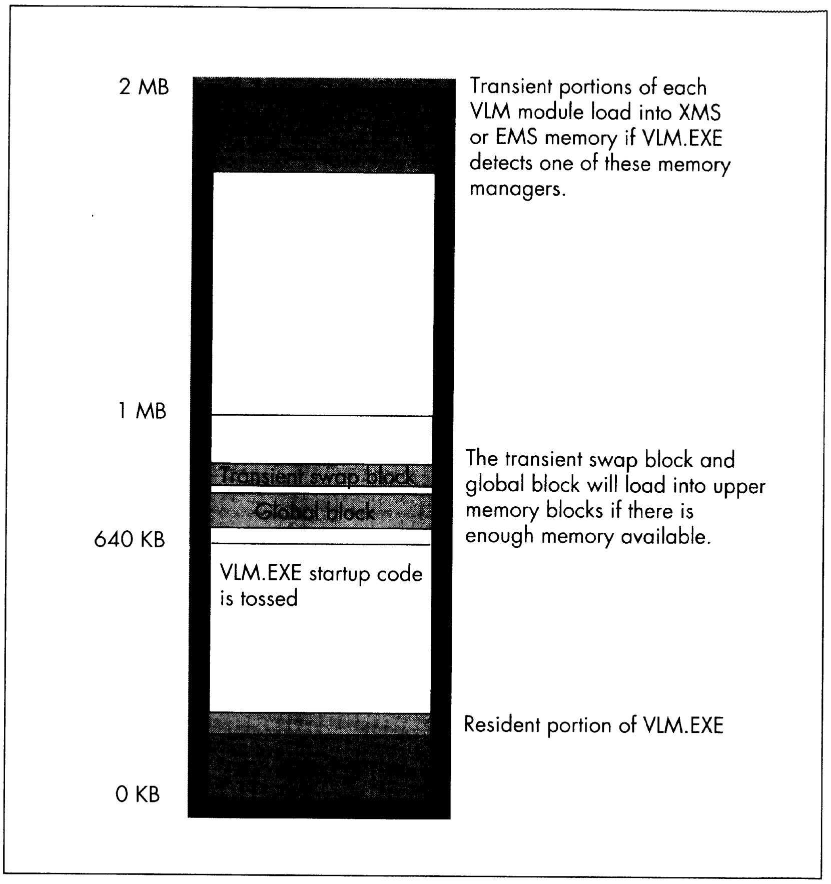
2 MB
Transient portions of each VLM module load into XMS or EMS memory if VLM.EXE detects one of these memory managers.
1 MB
Transient swap block
Global block
640 KB
The transient swap block and global block will load into upper memory blocks if there is enough memory available.
VLM.EXE startup code is tossed
Resident portion of VLM.EXE
0 KB

After all the VLMs pre-initialize, VLM.EXE begins the "real" initialization phase by calculating the total memory requirements for the *global blocks* and transient *swap block.* The global blocks, which are comprised of the global segments of each of the VLM modules you are loading, are roughly the size of the sum total of global segments. The transient swap block, during initialization, is equivalent in size to the largest VLM module being loaded. After initialization, the transient swap block is reduced in size to match that of the largest individual VLM transient segment.

VLM.EXE also determines which type of memory (XMS, EMS, or conventional memory) to allocate for storing the transient portion of the VLM modules. For example, once the DOS Requester pre-initializes, it knows how much memory to request from the XMS or EMS memory manager in order to store the transient portion of all the modules. If you don't have an XMS or EMS memory manager loaded, the transient portions of the modules are loaded into conventional memory, which can severely limit the size of the applications you can load.

The command line parameters `/mx`, `/me`, or `/mc` (for extended, expanded, or conventional memory) determine where the transient portions of the VLM modules are located. VLM.EXE uses this priority pattern for allocating memory when you don't specify a memory type or if the memory type you have specified is not available.

Each VLM module temporarily loads into the transient swap block and executes its initialization code, which includes initializing internal variables, hooking interrupts, notifying other VLMs of its presence, and detecting the presence of other dependent VLMs.

If the initialization of that module is successful, VLM.EXE splits the VLM module into its transient and global components, based on the memory requirements initially supplied by the VLM during the pre-initialization phase. The VLM's global and transient blocks are at this point relocated to their respective memory locations, with the global block being relocated to conventional memory or UMBs and the transient block being relocated to XMS, EMS, UMBs, or conventional memory, depending on the memory configuration. If all this is successful, the initialization code is jettisoned.

How the DOS Requester Performs a Read Request

Assuming that an application has successfully opened a file on the server and that the pointer within the file is positioned to where a read begins, let's examine how the NetWare DOS Requester performs a read request.

First, the application makes a call to DOS through its Int 21 service using the 3Fh "Read from file" function. Upon entry, the registers contain the function number (3Fh), the file handle that is passed back from the "Open File" request, the number of bytes to read, and an address to a buffer in memory to receive the data read from the file.

The Int 21 call is first intercepted by the NETX.VLM module (if loaded) to see if the call is within its range of serviced functions (B0h through E3h). Since Int 3Fh is outside this range, the NETX.VLM passes control to DOS.

DOS ascertains that the handle came from a remote device, which can be the network or a CD-ROM (DOS doesn't know which nor does it care). DOS generates a call using the Network Redirector Service Interrupt 2F to read from a remote device. DOS passes control to the REDIR.VLM module, which determines whether the read request is for the network. To do this, the REDIR.VLM module looks at the information stored in DOS's system file table. After determining that the request is for the network, REDIR sets itself up to handle possible record locking scenarios and passes control to the FIO.VLM module.

FIO.VLM receives the read call from REDIR.VLM via the VLM Manager. The FIO module gathers specific file information based on the settings used to open the file (is it read-only, is it sharable, etc.). At this time the FIO module determines whether the DOS Requester can cache reads, whether Packet Burst is enabled, and whether the call can effectively use Packet Burst. Packet Burst information is based on the packet frame size, the amount of information that FIO can pull from workstation cache (if a cache buffer is free to receive the read), and the size of the read request itself.

If the request can be completely fulfilled from data already found in the workstation's cache, the FIO.VLM module retrieves the data from the cache buffer and sends the data to the application without ever sending a packet across the wire. If the request can be partially fulfilled from data found in the workstation's cache, the FIO module retrieves that data and sends a request for the balance to the server through the IPXNCP module and the ODI stack.

If the workstation has caching enabled and there is an available cache buffer for the file (based on the NET.CFG settings and the type of file access being performed), the FIO module implements a read-ahead algorithm. This algorithm maximizes the efficiency of network read requests and cache, especially if the requested data is in relatively small increments.

Let's say an application has a request for a 2000-byte read and the cache buffer size is roughly 1450 bytes, which matches (roughly) Ethernet's 1514 frame size (minus 64 bytes of overhead). The first 1500 bytes of incoming data is handed to the application and is not cached. However, the second 1500-byte packet is cached, and the first 500 bytes are sent to the application. But the last 1000 bytes (approximately) of the second packet are the read-ahead for the next read request. If the next application request is within the 1000 bytes, it is fulfilled from the data stored in the cache buffer. If the request cannot be fulfilled from cache, the FIO module will clear the cache and perform the next read following the procedure explained above, which starts the process over again.

If the amount of data requested warrants using Packet Burst, the request is passed into the packet burst procedure of the FIO.VLM module and the request is bursted across the network. Packet Burst can greatly enhance the efficiency of read or write requests by using a one-request/multiple-reply method for data retrieval as opposed to the one-request/one-reply method that the NCP Read normally implements. If the workstation is not using Packet Burst, the workstation passes many more packets between itself and the server to fulfill the request. (For more information on Packet Burst technology, see the "Packet Burst Update: BNETX vs. VLM Implementation" AppNote in the November 1993 issue of *NetWare Application Notes.*)

After the FIO.VLM module completes the read, REDIR.VLM updates pertinent information in DOS's system file table, and the data and return codes are passed back through DOS to the application. The data requested is now in the application-supplied area of memory.

DOS Requester Performance Considerations

The DOS Requester and the NetWare Shell exhibit several different characteristics during the initial attaching and logging in process that tend to create a perceived speed gap between the two. Both the NetWare Shell and the DOS Requester query the network to find the nearest server for an initial attachment. They both obtain information from the initially attached server, and, if you specify a preferred server, they both obtain information about the preferred server and attach to it.

In addition, the DOS Requester executes security validations and performance enhancements that require this information to be sent over the network. These initial validations slow down the login process, but the performance enhancements greatly increase performance once you are logged in.

For example, the DOS Requester must perform NDS-specific functions during login that the NetWare Shell cannot, such as background authentication to multiple servers within the Directory tree, resolving context-specific information for NDS-style drive mappings, and so forth. NETX.EXE can only utilize bindery-based services on NetWare 4 servers.

On the performance aspects, the DOS Requester adjusts for Large Internet Packet support if you have LIP support enabled. You enable LIP support through the `Large Internet Packets = ON` (default) setting in the NET.CFG and if you have LIP support at the routing server. To negotiate LIP support, the DOS Requester performs an "LIP Echo" sequence to the destination server in order to establish the largest packet size that can pass through the router. This allows the DOS Requester to utilize a larger packet size for transferring data and thereby improve performance. If the NetWare Shell detects that the preferred server is more than one hop away (going through one or more routers), the packet size is forced to 512 bytes to compensate for the possibility of smaller frame sizes along the route.

The DOS Requester also negotiates a Packet Burst connection with the server if you have packet burst enabled through the `PB Buffers = 1` or more (anything other than 0) in the NET.CFG and if you have Packet Burst enabled at the preferred server. Packet burst processes large read requests much more efficiently and with much greater overall throughput, especially with WAN links. The NetWare Shell with Packet Burst enabled (BNETX.EXE) supports packet burst but in a limited fashion (and uses more memory than the regular NETX.EXE shell).

The DOS Requester manages both DOS's internal tables and the NETX.VLM module's internal tables in its effort to maintain full compatibility with the NetWare Shell. In some instances, this setup requires the DOS Requester to issue multiple requests in order to update the tables it uses, while the NetWare Shell issues just one request to maintain its internal tables.

You only need to load the NETX.VLM module with the DOS Requester if you want to supply backward compatibility with NetWare utilities that were written before NetWare 4 and with applications that are written to the pre-NetWare 4 specification. But in most cases, the NETX.VLM module needs to be loaded in order to handle the bulk of the NetWare-aware applications and utilities currently available on the market. If you are exclusively logging in to NetWare 4 servers through NDS LOGIN.EXE and are only using NDS utilities, you do not need to load the NETX.VLM module.

DOS Requester Performance Examples

Even given the above ramifications, the DOS Requester has had a checkered past to its overall speed and performance. With that in mind, here are three examples showing how to set the NET.CFG with different considerations.

Example 1 shows how to optimize the DOS Requester for the most speed but without memory considerations. Example 2 shows how to optimize a workstation's memory when running the DOS Requester. Example 3 tries to find a happy medium between Example 1 and Example 2.

Example 1: Optimization Considerations. The first NET.CFG configuration gives you the largest memory hit but offers the best performance. For the absolute best performance, use the `/MC` parameter when loading VLM.EXE, such as `VLM /MC <Enter>`. The `/MC` parameter loads all the VLM modules into conventional memory, which amounts to about 100 KB. (If you have to do this, load everything else that you can into upper memory so the memory hit isn't so bad.)

While all performance parameters are shown here, those parameters marked with an asterisk (*) do not need to be set because their default settings are already set for maximum performance. However, they are shown here for the sake of clarity. Parameters not listed here do not adversely affect performance.

In the following listing, as in examples 2 and 3, all lines would be indented under "NetWare DOS Requester."

```
 CACHE BUFFERS       = <FILES= in CONFIG.SYS - 5, or 64>
*CACHE BUFFER SIZE   = <max media size of NIC - 64>
*CACHE WRITES        = ON
 CHECKSUM            = 0 (only for use with 802.2 clients)
```

```
*LARGE INTERNET PACKETS     = ON
*LOAD LOW CONN              = ON
*LOAD LOW IPXNCP            = ON
*MINIMUM TIME TO NET        = 0
 NETWARE PROTOCOL           = NDS,BIND,PNW
*PB BUFFERS                 = 3 (any amount from 1 to 9)
 PRINT BUFFER SIZE          = 256
 SIGNATURE LEVEL            = 0
*TRUE COMMIT                = OFF
```

Setting Cache Buffers to match the `Files=` entry in the CONFIG.SYS (minus 5) gives almost a cache buffer for every file that you may have open (the maximum setting is 64). But what you set this to depends on if you are running Windows off the network or if you have a lot of network applications and files open concurrently.

It also depends on how much memory you have in conventional memory to give to cache buffers, because cache buffers are based on the maximum media size of the packets the workstation is sending. So if you're using Ethernet packets set to 1514 bytes, then each cache buffer will be about 1450 bytes (minus 64 bytes). If you set Cache Buffers to 64, you stand to tie up about 90 KB of conventional memory for cache buffers, control blocks, and other miscellaneous DOS Requester overhead.

The maximum cache available is 64 KB and is calculated by multiplying Cache Buffer Size by Cache Buffers. The FIO.VLM adjusts your cache buffers accordingly. If you set the number of Cache Buffers too high, you receive a warning (if your Message Level parameter in the NET.CFG is set to 3 or higher) telling you that the FIO.VLM is reducing your cache blocks by an amount. This is necessary for the FIO to load.

Note: Each VLM can be a maximum of 64 KB for the combination of the transient and global areas. Cache buffers are allocated after the other FIO.VLM requirements are met. Therefore, maximum cache buffer memory available equals 64 KB minus the other FIO.VLM allocations.

The `Cache Writes = ON` parameter fills the local cache buffers before writing data to the network, so using this default value is a good idea. If you don't need the added protocol checking, set the Checksum setting to 0 (it defaults to 1). For greatest performance, the next four settings, `Large Internet Packets` through `Minimum Time To Net`, should all be left to their default settings.

For the NetWare Protocol entry, put the server type that you use the most as first in the list. For example, set `Netware Protocol = BIND,NDS` if you primarily use 3.x servers, or `NDS,BIND,PNW` if you wish to log in under Directory Services first. To be set to ON, the `PB Buffers=` entry should be set to a number from 1 to 9 (OFF = 0). For best performance, be sure to have this turned ON.

Setting the `Print Buffers Size` to its maximum of 256 bytes gives you the best performance for printing. The `Signature Level` parameter offers NCP packet signatures, which can affect the overall performance of every packet a workstation receives and sends. For maximum performance, set `Signature Level = 0` in the NET.CFG and at the server, type `SET NCP Packet Signature Option = 0` to disable server packet signatures.

The `True Commit` entry increases data integrity but decreases performance because you have to wait for data to be written to the server's disk, not simply to the server's cache. For best performance, use the default of OFF.

Example 2: Memory Considerations. The memory configuration example uses the absolute least amount of memory for running the VLMs. However, please note that by doing this, you will see performance degradation.

For the best use of memory, have XMS memory loaded (for MS-DOS, load HIMEM.SYS in the CONFIG.SYS) and have a large contiguous block of UMBs (upper memory blocks) available (load EMM386.EXE in the CONFIG.SYS).

One way to allocate more UMBs is to disable expanded memory, which eliminates the allocation of an EMS page frame, saving you 64 KB of memory in the UMB area. (To do this, use the `/NOEMS` switch with MS-DOS's EMM386.EXE.)

Also be sure to have the lines `DOS=HIGH` and `DOS=UMB` in the CONFIG.SYS so MS-DOS will load into HMA (High Memory Area), thereby freeing up more conventional memory. The `DOS=UMB` parameter enables the UMBs for loading TSRs and other things, thereby freeing up more conventional memory as well.

Again, those parameters preceded with an asterisk (*) don't need to be set because their default settings are already set for maximum performance.

```
*AUTO LARGE TABLE       = OFF
 AUTO RECONNECT         = OFF
 CACHE BUFFERS          = 0 (performance hit)
 AVERAGE NAME LENGTH    = <calculated value, 48—default>
 CONNECTIONS            = <calculated value, 8—default>
 LOAD LOW CONN          = OFF (performance hit)
 LOAD LOW IPXNCP        = OFF (performance hit)
 EXCLUDE VLM            = <vlm>
 NETWORK PRINTERS       = 0
 PB BUFFERS             = 0 (performance hit)
 PRINT HEADER           = 0
 PRINT TAIL             = 0
 SIGNATURE LEVEL        = 0
```

As you can see, while many of these settings minimize the DOS Requester's memory requirements, they come with substantial performance hits.

The `Average Name Length` can be trimmed down by taking the sum of the lengths of all the server names you attach to, dividing by the number of servers you attach to, and then adding one. This entry works in conjunction with the `Connections=` parameter.

The `Connections` parameter shows the number of servers you attach to. The default is eight, but if you're using NDS, you may need to increase this number in order to handle "tree walking" through large corporate trees.

Use the `Exclude VLM = <vlm>` to eliminate optional VLMs that you don't need, which may include AUTO.VLM, NDS.VLM, BIND.VLM, or PNW.VLM. Some VLM modules may be optional to your configuration and do not need to be loaded. There are several ways to *not* load an optional module, including the following:

- Include the line `EXCLUDE VLM = <vlm>` under the "NetWare DOS Requester" heading in the NET.CFG file. This is the recommended way to disable a module.
- Rename the module with a different extension (i.e., .SAV).
- Delete the module (not recommended).

Optional VLMs include the following:

AUTO.VLM
PRINT.VLM (unless you want to print)
SECURITY.VLM
NDS.VLM (unless you are running Directory Services)
BIND.VLM (unless you connect to 3.1x servers and below)
PNW.VLM (unless you connect to Personal NetWare)
NETX.VLM
RSA.VLM
WSSNMP.VLM*
WSREG.VLM*
WSASN1.VLM*
WSTRAP.VLM*
MIB2IF.VLM*
MIB2PROT.VLM*
NMR.VLM*

(Modules marked with an asterisk are for SNMP services and will be covered in a future AppNote.)

Set the `Network Printer` to 0 only if you want to disable printing. The PRINT.VLM module will not load in this case. For the `Print Header` and `Print Tail` entries, you can set them to 0 if you are using network printing but you don't use any print job definitions set up through the PRINTCON utility.

Example 3: The Best Compromise. The best compromise of memory and performance is to use the default settings for the NET.CFG parameters, except when circumstances warrant modification. When this happens, pull from Examples 1 and 2. Example 3 is a list of the default parameters that affect both memory and performance. Since these are the default settings, all settings are marked with an asterisk.

```
*AUTO LARGE TABLE           = OFF
*AUTO RECONNECT             = ON
*AVERAGE NAME LENGTH        = 48
*CACHE BUFFERS              = 5
*CACHE BUFFER SIZE          = <max media size of NIC - 64>
*CACHE WRITES               = ON
*CHECKSUM                   = 1
*CONNECTIONS                = 8
*LARGE INTERNET PACKETS     = ON
*LOAD LOW CONN              = ON
*LOAD LOW IPXNCP            = ON
*MINIMUM TIME TO NET        = 0
*NETWARE PROTOCOL           = NDS BIND PNW
*NETWORK PRINTERS           = 3
*PB BUFFERS                 = 3
*PRINT BUFFER SIZE          = 64
*PRINT HEADER               = 64
*PRINT TAIL                 = 16
*SIGNATURE LEVEL            = 1
*TRUE COMMIT                = OFF
```

Conclusion

This Application Note details much of the pertinent theory on how the DOS Requester 1.1 works in a client environment. In particular, we covered the DOS Requester's theory of operations and how the NetWare Shell and the DOS Requester work with DOS. We also covered how the DOS Requester initializes, then how the DOS Requester performs a read request.

Finally, we looked at three NET.CFG examples that show some different ways to configure the DOS Requester. The examples looked at performance optimization, memory optimization, and the default settings.

This AppNote will be followed by another AppNote that explains the new NET.CFG settings. We'll also look at the new NET.CFG settings that come with Personal NetWare. Finally, we'll look at some tips, tricks, and traps on running the DOS Requester that you should know about.

Compression and Suballocation in NetWare 4

Alan Mark
Corporate Integration Specialist
Systems Engineering Division

It often seems that no amount of server disk space is sufficient for the storage demands of today's networks. Dramatic increases in the size of both applications and data files have taken their toll on disk storage resources. NetWare 4 offers several features to help alleviate this problem. This AppNote disccuses two of these features—file compression and disk block suballocation. It explains how they work and presents some test results to give you an idea of how much money you can expect to save by using these features on your NetWare 4 servers.

Trademarks

NetWare, the N-Design, and Novell are registered trademarks and the NetWare Logotype (teeth), Internetwork Packet Exchange, IPX, and SPX are trademarks of Novell, Inc. UNIX is a registered trademark and UnixWare is a trademark of UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc.

Macintosh is a registered trademark of Apple Computers Inc. IBM and OS/2 are registered trademarks of International Business Machines Corporation. Microsoft and MS-DOS are registered trademarks of Microsoft Corporation. All other product names mentioned are trademarks of their respective companies or distributors.

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of these Application Notes (AppNotes) or of any of the third-party products discussed in the AppNotes. Novell reserves the right to revise these AppNotes and to make changes in their content at any time, without obligation to notify any person or entity of such revisions or changes. These AppNotes do not constitute an endorsement of the third-party product or products that were tested. Configuration(s) tested or described may or may not be the only available solution. Any test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state, or local requirements. Novell does not warranty products except as stated in applicable Novell product warranties or license agreements.

Novell, Inc.
122 East 1700 South
Provo, Utah 84606 USA

Contents

Acknowledgements

Special thanks to Kyle Unice and Key Furner of Novell's Software Engineering Department, and to Ed Liebing of Novell Systems Research, for their help with this AppNote.

Introduction

With the release of NetWare 4, Novell has introduced three new server technologies to more effectively utilize disk storage on your NetWare server:

- File compression
- Disk block suballocation
- Data migration

This AppNote explains the internal workings of the first two methods and discusses the impact of using these technologies on your servers.

Compression: An Overview

Even though the price of disk storage has dramatically declined over the past 10 years, the need for data compression has risen due to two main factors.

First, today's applications require more storage space than in the past. Remember the days when an entire application would fit on one floppy disk? Now some vendors (including Novell) are shipping their software on CD-ROMs for easier and faster installation.

The second reason that compression is needed today is for efficiency. It's doubtful that most of us really use all data that's available to us on file servers or on our local hard disks. So rather than having data always expanded and taking up valuable disk storage, it makes sense to have infrequently-used data be readily available, but in a compressed format.

Also, some file types are inefficient. For example, a simple, black-and-white Windows Paint file (BMP) I created took up 812,086 bytes of disk space. When compressed, the size went down to 1,948 bytes.

How Compression Works

The goal of compression algorithms is to rearrange or encode data in such a way that the resulting data is a fraction of the original data's size.

A common data compression method, called the *duplicate string encoding algorithm*, analyzes data for redundant patterns and then asssigns a code for each pattern. Depending on the implementation, the process might search the data by bits, bytes, or double-bytes. The resulting compressed file is actually an encoded file of nearly random data that has a list of patterns which, when pieced back together, form the original file. Data that is nearly random usually cannot be further compressed.

To illustrate how duplicate string encoding compression works, let's use the following famous quote and see how the resulting compressed file might look. Remember that this is a very basic example—real compression algorithms are much more complex.

"Ask not what your country can do for you; ask only what you can do for your country." (84 bytes)

Obviously, most files don't contain such repetative data, but using a basic encoding scheme might result in the following header for the coded file (the underscore [_] represents a space):

1=sk 2=_not 3=_what 4=_you 5=_country 6=_only 7=_can 8=_do 9=_for

The resulting data within the coded file would be:

A1234r56789; a16347894r5. (25 bytes)

So the size of the resulting encoded file would be 25 bytes plus the header (decoding) information. Compressing small files usually results in negligable savings due to header overhead. For example, using PKZIP on the 84-byte file above resulted in a 176-byte file—92 bytes *larger* than the original file. For this reason, NetWare 4 only compresses files larger than 512 bytes.

Better compression ratios are the result of finding as many patterns as possible. This may mean keeping extensive pointers to data already processed (these are known as "back pointers"). The way in which patterns are discovered directly affects the time it takes to process the data. Therefore, a happy medium must be found for background compression to take place on NetWare servers so that decompression times are quick while compressed size is small (see Figure 1).

There are many other compression algorithms in the computer world, and some work best on specific data types. For instance, certain algorithms are designed solely for video and sound; others are for scanned images. There also exists "lossy" compression routines which offer greater compression ratios but do so by losing extraneous data that won't be relevant when decompressed. These routines are only acceptable for video (where human interpretation makes up for lost information), pictures (where colors are reduced because the eye cannot discern them), and sound (when our ears will not discern background sounds). It's obvious that most computer data must be restored to their original state (called "lossless" compression).

Figure 1: NetWare server compression strikes a happy medium between high compression ratios and quick compression times.

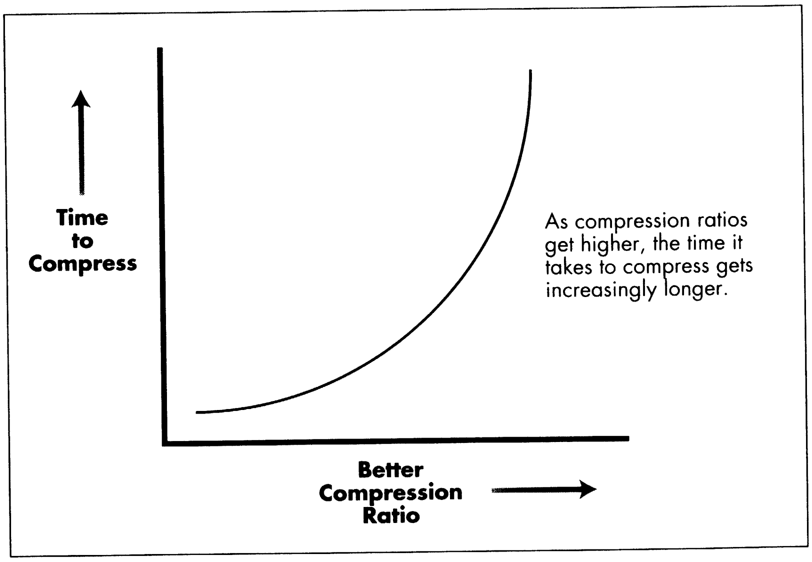

Common Compression Implementations

Before we get into NetWare 4 server compression, let's briefly discuss some other ways that compression is used in today's PC environments.

File Compression. Anyone who has dialed into a bulletin board system or online service such as CompuServe has worked with compressed files. Utilities such as PKZIP (from PKWare, Inc.) and Stuffit (from Aladdin Systems, Inc.) are used to compress individual files to save on download time or to place more data onto a floppy disk. Some files, such as those in CompuServe's GIF format, are by definition always compressed and require specialized utilities to access the data.

File compression utilities work by reading each file and compressing it into a new file (usually called an *archive* file). They also have the option of creating a self-extracting archive (SEA), so that the decompression code is built into the file. SEAs can then be decompressed by people who don't own the compression utility. However, SEAs are applications that only operate in their native environment (DOS, Macintosh, or Unix).

File compression utilities are great for online systems and for sending files via e-mail. But they require a lot of user intervention and seldom operate automatically. Thus they are used only in specific situations.

Disk Compression. Recent versions of MS-DOS and Novell DOS include disk compression programs which go by such names as DiskDoubler and Stacker. (Other software companies also offer similar programs.) These programs work by creating a new disk volume (also called a *virtual* disk) on an existing drive, and then compressing and moving all data over to the new compressed volume. From then on, all data on the volume is compressed and decompressed in real-time; you don't get to choose which files get compressed and which ones don't.

Compression is performed as data is written to or read from the volume without user intervention. This method of transparent data compression is extremely easy to use and gives the user about 50 percent more storage than before, depending upon the data stored.

For more information on disk compression, see "Exploring Hard Disk Compression" in the August 1993 *NetWare Application Notes.*

The main drawback to disk compression is that it slows down the system because every disk access requires a compression process to be executed. In recent industry tests of MS-DOS and PC-DOS disk compression, the results indicate that—depending upon the application—disk compression slowed performance from a little to over fifty percent (see *InfoWorld*, January 24, 1994.)

Compression in Communications. A more recent application of compression is in the area of data communications involving modems and LAN/WAN links. Here, both hardware and software work to increase throughput on various types of topologies and media.

The current wave of modems incorporate a chip that compresses data in real-time during an active communication session. Using these compression methods (which go by such names as MNP5 and V.42 *bis*), a 9600-baud link can send and receive data several times faster than without compression. However, each modem must use the same technology for compression to work.

For more information on compression over WAN links, see "Optimizing NetWare Wide Area Networks" in the May 1994 *Novell Application Notes.*

Another new trend is to enable software compression for remote communications in products such as NetWare Remote Node and NetWare Connect. Here, the remote user and local server use V.42 *bis* implemented at the network level to enable compression for all applications.

Compression is also being used by routers to increase throughput on a wide area network. For instance, with the NetWare Multiprotocol Router and appropriate hardware, network packets can be compressed.

NetWare 4 Server Compression

Novell engineers developed a new compression method specifically for NetWare 4. It allows the network manager to maximize disk storage using *intelligent* criteria. The other methods discussed above either work on *specified* files or on *all* files, but NetWare 4 allows you to better control the compression environment.

NetWare 4 server compression operates in two ways. The manual method is to mark a file or directory as "Immediate Compress" with the FLAG command (see section below). But the best way to handle file server compression is to let the server itself determine which files to compress and when. This is done by adjusting several SET commands on the server console.

NetWare compression works for all name spaces including Macintosh and Unix file systems, and therefore works seamlessly on all platforms. It is enabled by default when a NetWare volume is created, and can be enabled on an existing volume using the INSTALL utility. Keep in mind that once compression is installed on a volume, it cannot be removed. However the compression routines can be enabled and disabled using a SET parameter, as explained later.

Before we get into how server compression works, let me review the NetWare file attribute called "Last Accessed." Novell DOS and MS-DOS only keep track of when a file was last modified. The Macintosh file system also records file creation dates. But for intelligent file compression to occur, the operating system must keep track of the last time a file was *accessed* (read from or written to). With this piece of information, NetWare can compress files when they have not been touched after a specified length of time. (Unix also records last access time.)

How NetWare 4 Compression Works

In very simplified terms, the NetWare 4 compression algorithm works as described above: duplicate strings are encoded and the compressed file becomes a coded image of the original file.

Compression begins when one of the following conditions occurs:

1. The Immediate Compress (IC) bit is set on a file or directory.
2. A file has been deleted and the salvageable file system has been enabled for immediate compression.
3. The nightly search thread has discovered an untouched file or salvageable file and queued it for compression.

These three conditions are illustrated in Figure 2.

Figure 2: Events that cause files to be queued for compression.

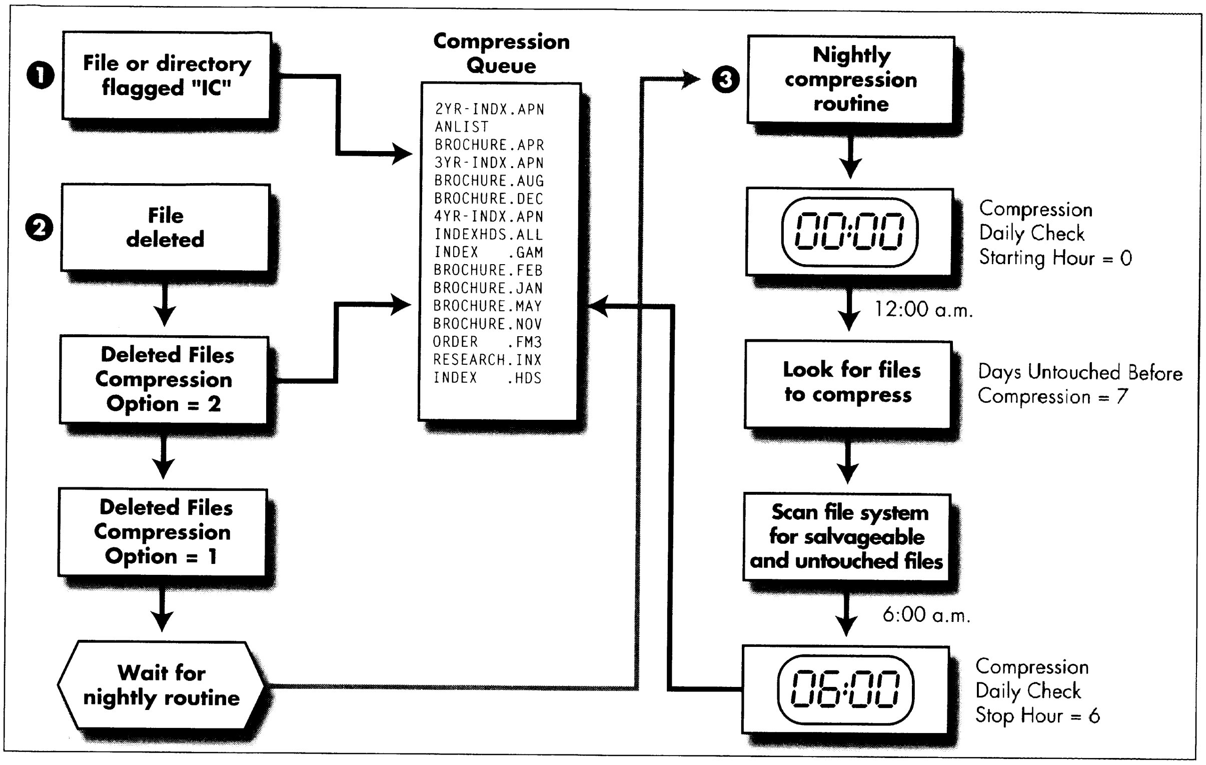

1
File or directory flagged "IC"
2
File deleted
Deleted Files Compression Option = 2
Deleted Files Compression Option = 1
Wait for nightly routine
Compression Queue
2YR-INDX.APN
ANLIST
BROCHURE.APR
3YR-INDX.APN
BROCHURE.AUG
BROCHURE.DEC
4YR-INDX.APN
INDEXHDS.ALL
INDEX .GAM
BROCHURE.FEB
BROCHURE.JAN
BROCHURE.MAY
BROCHURE.NOV
ORDER .FM3
RESEARCH.INX
INDEX .HDS
3
Nightly compression routine
00:00
Compression Daily Check Starting Hour = 0
12:00 a.m.
Look for files to compress
Days Untouched Before Compression = 7
Scan file system for salvageable and untouched files
6:00 a.m.
06:00
Compression Daily Check Stop Hour = 6

For you techies out there, the compression thread runs at low priority. After processing 128 bytes of the original or temporary file, the thread reschedules itself with the RescheduleLastLowPriority call.

As a file is queued for compression, the original file is analyzed and a temporary file is built with translation information. This temporary file is kept in the disk cache by the caching subsystem, as long as the temporary file doesn't consume more than half of the cache buffers.

Although the compression routines have been optimized for size, they are queued as secondary processes so that ordinary disk requests will be executed first.

After the temporary file has been built and the file size for the compressed version has been calculated, NetWare determines whether any disk sectors will be saved by having compressed the file. The value of the SET parameter "Minimum Compression Percentage Gain" (default is 2%) is compared to the calculated savings. If the savings is greater than the "Minimum Compression Percentage Gain," and no errors have occurred during the process, the compression thread will begin creating the compressed version of the file by processing the temporary file.

Only after the compressed file has been completely written will a controlled swap of the original and newly compressed file take place. If any errors—including power failure—occur during this process, the original file is left intact.

Compression routines are CPU intensive and thus are best left to execute during off-hours. This is why the default for the "Compression Daily Check Starting Hour" parameter is set to 0 (midnight) and the default for the "Compression Daily Check Stop Hour" parameter is set to 6 (6 a.m.). The related parameter is "Days Untouched Before Compression," which has a default of 7 days. Using these defaults, the server will scan the file system from 12 a.m. to 6 a.m. and look for files left unaccessed for at least seven days before it schedules compression on those files.

If a compressed file is accessed, one of two things happens after the file has been decompressed—depending on the setting of "Convert Compressed to Uncompressed Option" parameter:

- If set to 0, the file will remain on the disk as compressed.
- If set to 1, the file will remain compressed, provided that it is not again accessed before the time specified by the "Days Untouched Before Compression" parameter.

For example, suppose that "Convert Compressed to Uncompressed Option" is set to 1 and "Days Untouched Before Compression" is set to 5 days. If I access a compressed file on Monday, the file will be decompressed into memory and sent to me, but will remain compressed on the disk. If I access the file again on Thursday, it will be permanently decompressed to disk as it is being retrieved and sent to me. If I had waited until Saturday to access the file, the file would still remain compressed.

Another intelligent use of server compression is found with "Deleted Files Compression Option." This parameter determines when and if deleted files should be compressed. If set to 1, then the deleted files will be compressed the next day (between the start and stop compression hours). If set to 2, the files will be compressed immediately.

The chart in Figure 3 summarizes the compression options settable via the SET command on the server. If you need to change any of the defaults, put the appropriate commands in the server's AUTOEXEC.NCF file so they will take effect whenever the server is booted.

Figure 3: SET parameters for file compression on a NetWare 4 server.

SET Parameter	Explanation	Default	Notes
Enable File Compression	Set to ON to allow compression on compression-enabled volumes. When set to OFF, compression will pause.	ON	When OFF, files flagged IC are queued until compression is allowed.
Minimum Compression Percentage Gain	If the compressed file won't be this much smaller than the uncompressed size, it is not compressed.	2 (%)	
Compression Daily Check Starting Hour	Specifies when to start the search for files that have not been accessed.	0	Hours are specified in military time (0=midnight).
Compression Daily Check Stop Hour	Specifies when to stop the search for files that have not been accessed.	6	Hours are specified in military time (0=midnight).
Days Untouched Before Compression	Specifies how many days a file must remain unaccessed before it can be queued for compression.	7	
Convert Compressed to Uncompressed Option	Specifies how the server stores a compressed file after uncompressing it. 0=always leave the file compressed; 1=leave it compressed after a single access within the "untouched" period; 2=always leave it uncompressed.	1	
Deleted Files Compression Option	Specifies how the server handles unpurged deleted files. 0=don't compress; 1=compress during the next search interval; 2=compress immediately	1	
Maximum Concurrent Compressions	Specifies how many compression operations can be performed simultaneously. Concurrent compressions can occur only on multiple volumes.	2	Increasing this setting may slow server performance.
Uncompress Percent Disk Space Free to Allow Commit	As a disk volume runs out of space, a compressed file which will be permanently decompressed may require too much disk storage. This parameter prevents newly decompressed files from using too much valuable free space.	10 (%)	
Uncompress Free Space Warning Interval	If files cannot be decompressed due to lack of free disk space, a warning is sent on the console. This parameter determines how often the warning is sent.	31 min. 18.5 sec	To prevent unnecessary temporary decompressions due to lack of disk space, the network administrator should monitor the server's free disk space to ensure that it exceeds the "Uncompress Percent Disk Space Free to Allow Commit" setting.

File and Directory Attributes Related to Compression

The following attributes are set using the FLAG utility or from other NetWare utilities.

Immediate Compress (IC). Changeable by client for files and directories. When set for a file, indicates that the file should be queued for compression. When set for a directory, indicates that all files placed into the directory should be queued for compression unless the DC or CC attributes are set for the file (see below).

If a file is copied into a directory flagged as IC, the file's attributes won't change. If a compressed file in a IC-flagged directory is moved to another directory, the file will remain compressed.

Don't Compress (DC). Changeable by client for files and directories. This attribute indicates that the file or directory should never be compressed. If set on a file that is already compressed, then the file will remain compressed until accessed. So setting a compressed file to DC won't immediately decompress it, but setting an uncompressed file to IC *will* immediately compress it.

If set on a directory, then all files in that directory will not be compressed unless a file is flagged IC. A file's attributes will not be changed if it is copied into a DC-flagged directory.

Can't Compress (CC). Set by the NetWare file system and unchangeable by the client. Indicates that the server has tried unsuccessfully to compress the file. Once this bit is set, it will not be reset until the file is written to. This attribute indicates that the file is less than 512 bytes in length, that the file is already compressed by another compression algorithm, or that the data within the file is nearly random.

File Compressed (CO). Set by the NetWare file system and unchangeable by the client. Indicates that NetWare has successfully compressed the file.

How Decompression Happens on the Server

File decompression is initiated anytime a file is opened, except when it is to be specifically accessed in its compressed state (see "NetWare Compression Tips" below).

The server reads the compressed file in 4KB chunks. As each chunk is retrieved, it is verified and decompressed into memory. The decompression routine notifies NetWare when 4KB of data is ready to be transmitted to the client. This has the effect of sending data as quickly as possible rather than having the client wait for the entire file to be decompressed.

If a file is set for "Immediate Compress" and is opened for reads or writes, the file is decompressed to disk. After the file is closed, it is queued for immediate compression.

For techies: The decompression thread relinquishes the CPU when 256 bytes are read from the compressed file, and when 4KB of decompressed data has been generated. The decompression work-to-do thread and its associated read-ahead thread both run at normal priority.

Note: Keep in mind that, when saving data, some applications create new files instead of updating existing files. In that situation, the new file will not have the IC attribute. Thus it is best to flag the directory, rather than individual files, to be compressed.

Compression Ratios

The following table shows typical compression ratios that can occur on a NetWare 4 server.

Figure 4: Typical NetWare 4 compression ratios.

File Type	Original Size (bytes)	Compressed Size (bytes)	Compression Ratio
Bitmap image (BMP)	32,078	15,360	52%
Windows DLL	451,280	227,840	50%
Text file	43,524	16,384	62%
Text file, table	111,947	22,528	80%
Windows executable (EXE)	1,039,904	514,560	50%

How does NetWare compression stack up against other popular file compression utilities? The following table gives you a general idea of how much the most popular utilities compressed the five various files listed above (BMP, DLL, text, text table, and Windows EXE).

Figure 5: Comparison of compression ratios of five file types by various utilities.

Compression Utility	Compression Ratio (%)
NetWare 4	53%
PKZIP 2.04g (DOS)	56%
Stuffit Deluxe 2.01 (Macintosh)	51%
Compress (Unix)	40%

The time it takes to complete any compression routine is directly dependent upon the speed of the CPU and whether the routine is performed in the foreground or background. Therefore, I have not included the time it took to complete the routine.

Note that while NetWare lets all other disk requests supercede files being *compressed*, it must be able to quickly *decompress* files in real-time while performing other network tasks and without user intervention. So NetWare reads compressed files faster and easier than any of these compression utilities.

NetWare Compression Tips

Following are some tips for effectively using NetWare 4 compression.

- To copy compressed files in their compressed format, use NCOPY with the `/R` or `/RU` parameters.

  The `/R` parameter will keep files compressed only if the destination volume supports compression. This speeds up the copying of already compressed files.

  The `/RU` parameter will keep a NetWare-compressed file in that state even if the destination volume (such as your local hard disk) doesn't support NetWare compression. Be aware that copying NetWare-compressed files to a non-NetWare volume means that the files cannot be decompressed unless they are first copied back to a NetWare volume which supports compression.

- The NetWare 4 version of NDIR has a `/COMPressed` parameter that displays file compression information. A sample screen is shown below:

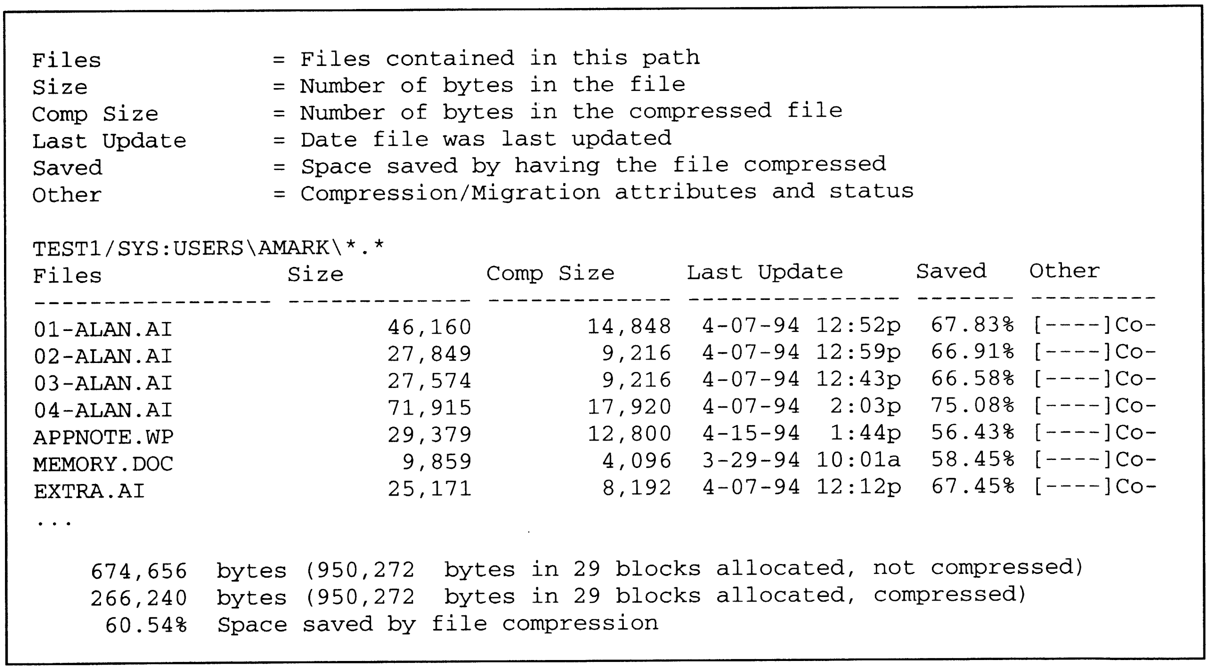

```
Files              = Files contained in this path
Size               = Number of bytes in the file
Comp Size          = Number of bytes in the compressed file
Last Update        = Date file was last updated
Saved              = Space saved by having the file compressed
Other              = Compression/Migration attributes and status

TEST1/SYS:USERS\AMARK\*.*
Files                Size            Comp Size     Last Update      Saved   Other
----------------- ------------- ------------- ---------------- ------- ---------
01-ALAN.AI                46,160        14,848  4-07-94 12:52p  67.83% [----]Co-
02-ALAN.AI                27,849         9,216  4-07-94 12:59p  66.91% [----]Co-
03-ALAN.AI                27,574         9,216  4-07-94 12:43p  66.58% [----]Co-
04-ALAN.AI                71,915        17,920  4-07-94  2:03p  75.08% [----]Co-
APPNOTE.WP                29,379        12,800  4-15-94  1:44p  56.43% [----]Co-
MEMORY.DOC                 9,859         4,096  3-29-94 10:01a  58.45% [----]Co-
EXTRA.AI                  25,171         8,192  4-07-94 12:12p  67.45% [----]Co-
...

    674,656  bytes (950,272  bytes in 29 blocks allocated, not compressed)
    266,240  bytes (950,272  bytes in 29 blocks allocated, compressed)
     60.54%  Space saved by file compression
```

- The `NDIR /VOL` command also displays compression information, as shown below:

```
Statistics for fixed volume TEST1/APPS:
Space statistics are in KB (1024 bytes).

Total volume space:                          512,000    100.00%
Space used by 437 entries:                    75,712     14.79%
Deleted space not yet purgeable:                   0      0.00%
                                            --------   --------
Space remaining on volume:                   436,288     85.21%
Space available to AMARK:                    436,288     85.21%

Maximum directory entries:                     8,704
Available directory entries:                   2,587     29.72%

Space used if files were not compressed:     194,189
Space used by compressed files:               82,770
                                            --------
Space saved by compressing files:            111,419     57.38%

Uncompressed space used:                      18,094
```

- An undocumented compression screen can be enabled on the server. Although the screen is designed for debugging purposes, it nonetheless lets the administrator see which files are currently being compressed. The server console command is SET COMPRESS SCREEN=ON, which displays information similar to the following sample:

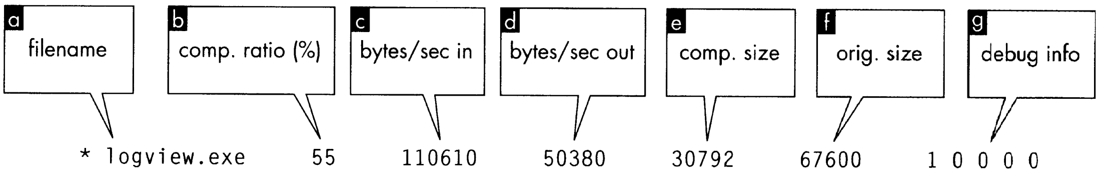

```
* logview.exe    55    110610    50380    30792    67600    1 0 0 0 0
```

The screen shows (a) the file being processed (an asterisk means decompression); (b) the compression ratio; (c) bytes/second processed into the compression engine; (d) bytes/second processed out of the compression engine; (e) the file's compressed size; (f) the file's original size; and (g) debug information having to do with how the file was processed, whether or not it could be compressed, and other codes useful to Novell programmers.

- There are significant issues concerning the backup of NetWare compressed files. If the backup program isn't NetWare 4-aware, each compressed file may be first decompressed before being sent to the backup device. If that same program then restores those files, they will be restored as uncompressed and require much more storage than before. Also, most backup programs restore the last accessed date to the original value.

  If the backup program *is* NetWare 4-aware and uses Novell's Storage Management Services (SMS), then compressed files will be backed up and restored as compressed.

Disk Block Suballocation

Many of you may be familiar with the term "cluster size," which refers to the minimum file allocation unit for local hard disks. On NetWare servers, the similar term is "disk block size."

NetWare 4 allows disk block sizes to be set to 4KB, 8KB, 16KB, 32KB or 64KB. Using the 64KB block size results in very fast disk operations. The reason is simple: The larger the disk block, the more data which can be transferred to and from the disk in a single request. So a 100KB file using 64KB blocks can be read in two disk reads, whereas it takes 25 disk reads with a 4KB block size. Also, larger blocks require less memory and the read-ahead operation operates faster. (Read ahead is a background task that reads sequential files into cache in advance of the request.)

However, there is a major disadvantage of using a large disk block size: Since at least one disk block is allocated per file, files that don't completely fit into a disk block will leave much unused space. With 64KB blocks, that can be a lot of wasted storage.

For instance, suppose that the disk block size is 4KB (4,096 bytes). Any file that is not an exact multiple of 4KB will have unused space in the amount of:

```
File size MOD 4096
```

(The MOD function returns the remainder of the division of the two numbers). So, for example, a 3KB file will have 1KB wasted, and a 4097-byte file will have 4095 bytes wasted. This methodology can result in the inability to completely utilize disk space. In other words, even though storage space is available, it cannot be used.

Figure 6 shows how files are stored in a disk block. This scheme works the same for DOS, Macintosh, and Unix systems and for NetWare servers without suballocation.

Figure 6: Traditional file storage schemes waste disk space.

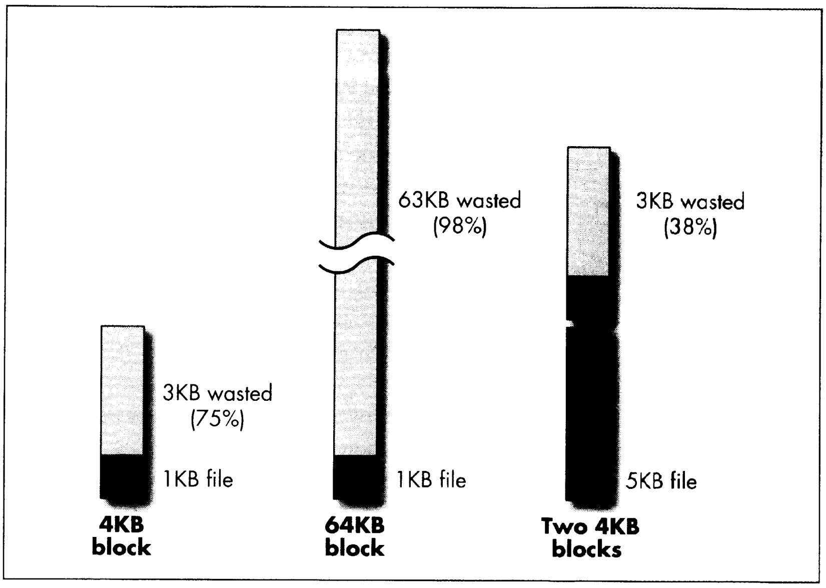

As shown in Figure 6, smaller files usually create the most wasted space as a percentage of their size. For instance, both 1KB and 5KB files have 3KB of wasted space, but the percentage is larger for the 1KB file (75% versus 38%).

Our discussion of suballocation will revolve around the following scenario: A NetWare 3.12 disk partition was divided into three volumes, each 23,000,000 bytes in size with a 64KB block size. (The same results would be reached if using NetWare 3.11.) The server was then upgraded to NetWare 4 and suballocation was enabled.

The Directory Entry Table and Its Role in Space Utilization

The maximum number of files which can be stored on a volume depends on the number of directory entries and blocks available. In NetWare 3.x and 4.x, the initial size of the directory entry table is determined according to the volume size, and then expands as needed. On our 23MB volumes, 512 directory entries were created.

Note: Remember that directory entries are used by subdirectories as well as files, and the directory entry table is itself a file.

Since each volume had a 64KB block size, the maximum number of disk blocks was 350, of which 347 were available (approximately 23MB / 64KB). (The precise calculation includes accounting for hidden files used by NetWare.)

Since each file requires at least one disk block, the most files that can be stored on a volume *without suballocation* is determined by the number of available blocks. Therefore, without suballocation, at most 347 files can be created on our 23MB volume. But with suballocation enabled, as we will see, significantly more files can be created because each file doesn't necessarily use one complete block.

Life Without Suballocation

How many files did it take to fill our NetWare 3.12 volumes?

On SYS2, I created a subdirectory and proceeded to fill it with 1KB files until the disk was filled. After the 347th file was created, a "Disk Full" message appeared on the screen.

On SYS3, I followed the same procedure, but this time with 32KB files. Once again, the 347th file caused a "Disk Full" error.

On SYS4, I created two directories and stored fifty 1KB files in \DIR1 and fifty 32KB files in \DIR2.

The results of these tests are shown graphically in Figure 7.

Figure 7: Amount of data that fit on the NetWare 3.12 volumes (without suballocation).

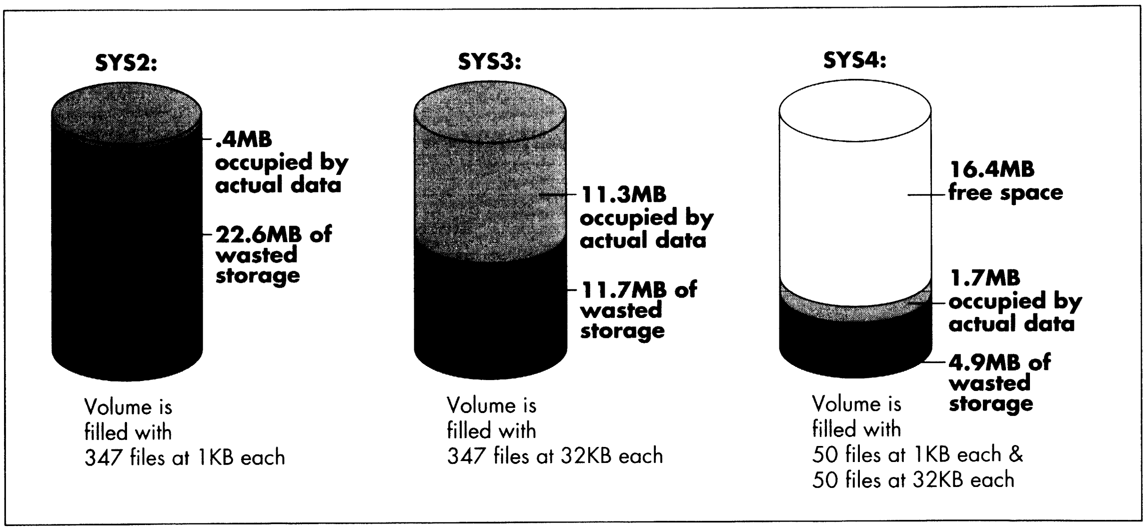

The numbers behind this graphic are given below. Notice the slack for each volume. Slack is the percentage of wasted space that could otherwise be used for storage. For SYS2, 98 percent of the disk space was wasted because each 1KB file took up the space of a 64KB file.

Volume	Block Size	Dir. Ent. Used	Data Stored	Storage Allocated	Wasted Storage	Slack
SYS2	64KB	347	347.0KB	23.0MB	22.6MB	98%
SYS3	64KB	347	11.3MB	11.7MB	11.7MB	50%
SYS4	64KB	102	1.7MB	6.6MB	4.9MB	74%

How NetWare 4's Suballocation Works

NetWare 4's suballocation feature is designed to solve the dilemma of how to increase server performance by using large disk block sizes, but not be penalized by the wasted storage that large blocks cause. By dividing a block into 512-byte chucks, the most wasted space you'll ever have is 511 bytes—even with 64KB blocks.

The suballocation routines in NetWare 4 are quite complex, but neither users nor administrators need to worry about it. There are no utilities or console commands to contend with, except to merely enable suballocation from the INSTALL program.

Files are traditionally controlled using the File Allocation Table (FAT), and in all versions of NetWare the FAT is hashed and indexed for faster data retrieval. NetWare 4 adds suballocation. When suballocation is enabled on a volume, the server designates blocks for a specific range of file sizes or ending data fragments (that is, data that cannot nearly or completely fill a block). Each "Suballocation Reserved Block," or SRB, is specific to a narrow range of file sizes or ending data fragments based on multiples of 512 (1 to 512 bytes, 513 to 1024 bytes, etc.).

Figure 8 illustrates this concept.

Figure 8: Suballocation reserved blocks.

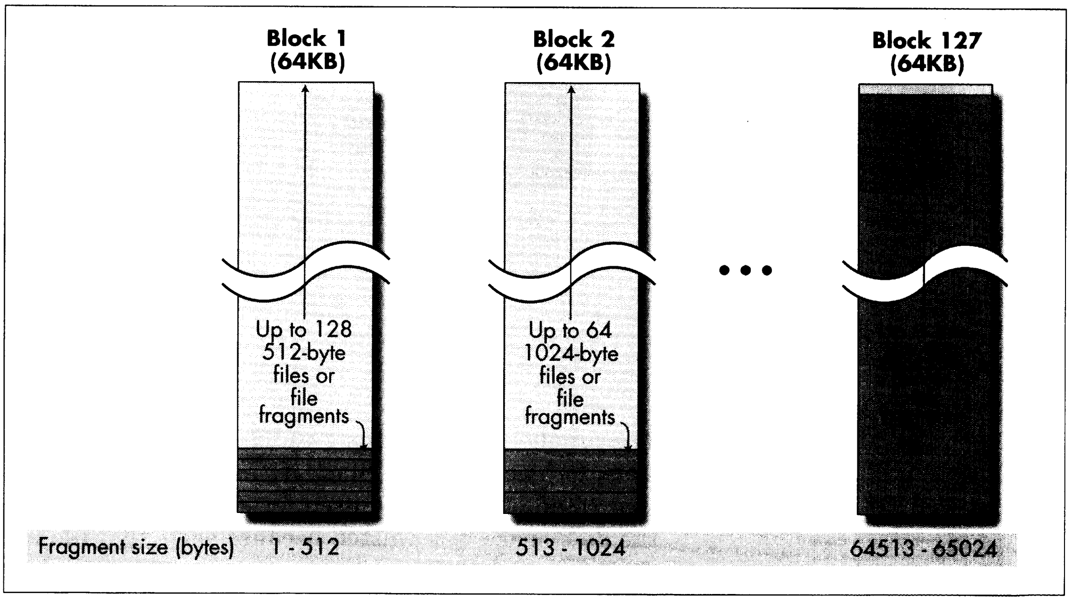

The number of SRBs is determined by the formula:

```
SRB = (block size / 512) - 1
```

For a 64KB block size, at most 127 blocks are reserved. Files that are within 511 bytes of the block size (e.g. from 65,025 to 65,536 bytes for 64KB blocks) are not suballocated.

So for 64KB blocks, the SRB designated to hold 1KB files can hold up to 64 files or ending data fragments, while the 20KB SRB can hold just over 3 files or fragments.

Once an SRB is full, another block is dynamically created and chained for suballocation use. Hence there can be chains of blocks that are used specifically for suballocation. Also, a file or fragment will be stored in two SRBs if the data cannot be completely stored within one SRB. For example, after a 20KB SRB contains three files or fragments, 4KB of the fourth file is stored in the current SRB and the remaining 16KB is stored in another SRB. This method optimizes the use of suballocation reserved blocks.

How does a file get suballocated? First, the server computes the fragment data size with the formula:

```
Fragment data size = file size MOD block size
```

For example, with a 64KB block, a 65,885-byte file will have a fragment data size of 65,885 MOD 65,536 = 349. If the fragment is within 512 bytes of the block size, the file is not suballocated.

Second, the server determines if an SRB exists in the range of the fragmented data. If not, it reserves a block. If one does exist but is filled or will cause an overflow, then a new SRB is created and the chain is extended. For a 65,885-byte file, the SRB used is for range from 1 to 512.

Finally, if the file is larger than one disk block, it is written into two sections of the disk: the primary disk block and the SRB.

All files with fragments in the same range will have their ending data stored in the same SRB chain. So a 65,885-byte file and a 65,886-byte file will have their ending data stored in the same chain because their fragments (349 and 350 bytes, respectively) fall within the same 512-byte range.

As files are created, expanded or contracted, data may be moved from one SRB to another. This may leave holes within the SRB. Since SRBs take up disk space, they are periodically cleaned up by compacting the entire chain and releasing unused blocks for primary storage. For instance, if a file is originally 4KB, and is then opened and closed so that the new size is 7KB, the updated file will be stored in an SRB designated for 7KB entries, leaving an empty space in the old 4KB-designated SRB.

Test Results with Suballocation

Let's return to our test server scenario. I performed the upgrade to NetWare 4 from the LAN; the process took about 20 minutes. I enabled suballocation for each volume, and then logged into the server and analyzed the volumes.

Just as I expected—no extra space was available. "What?" you say, "Even with suballocation turned on?" The reason is as follows:

> *For existing volumes, suballocation only affects newly-created files.*

This is an important point. If you upgrade existing NetWare 3.x volumes to NetWare 4, the benefits of suballocation won't be realized until existing files are opened or new files are written to the disk. Untouched files won't benefit.

Most NetWare 3.x volumes use small block sizes. Since suballocation works best for block sizes greater than 4KB, to fully get the benefits of large block sizes and suballocation, you should create new volumes with 64KB block sizes and suballocation enabled. So to implement this new technology on your newly upgraded NetWare 4 server, you'll have to create new NetWare volumes. There are two ways to do this.

For a detailed discussion of NetWare migration strategies, see "NetWare Migration Utilities Part 1: The In-Place Upgrade NLM" in the June 1993 issue and "NetWare Migration Utilities Part 2: The Across-the-Wire Migration Utility" in the September 1993 *NetWare Application Notes.*

The first method is to use over-the-wire migration, where a NetWare 3.x server sends data to a new NetWare 4 server over a LAN connection. The new server should have compression and suballocation turned on before the migration process begins.

The second method is to back up the NetWare 3.x server (onto tape or external hard disk), upgrade the same server to NetWare 4, create new NetWare 4 volumes, and finally restore the data to the new volumes.

For my server, I simply deleted all files on the volumes and re-ran the test up to the maximums reached in the first test. The results, as shown below, demonstrate the significant storage savings using suballocation *in a best case scenario* (one with no slack).

Volume	No.of File	Dir. Ent. Used	Data Stored	Storage Allocated	Wasted Storage	Slack
SYS2	347	347	347.0KB	347.0KB	0KB	0%
SYS3	347	347	11.3MB	11.3MB	0MB	0%
SYS4	100	102	1.7MB	1.7MB	0MB	0%

Of course, no real-life situation will give you 0 percent slack. But under normal use you can expect significant savings, especially on volumes which store many small files (even a 1-byte file is suballocated). As files are stored on a volume with a large disk block size, the more suballocation saves disk space. Also, suballocation works for files in other name spaces and on deleted files (deleted files continue to use disk space until purged).

To further test the effectiveness of suballocation, I copied a 19MB file to volume SYS2 and it fit! Remember that without suballocation, I couldn't copy *any* more files to SYS2.

Savings with Suballocation

How much more disk space can you expect with suballocation? I developed a spreadsheet to calculate how much suballocation saved us on our production server. The calculations are based on formulas developed by Novell engineers to determine server memory requirements (see Figure 9).

Figure 9: Spreadsheet for comparing cost of RAM and wasted disk space with and without suballocation.

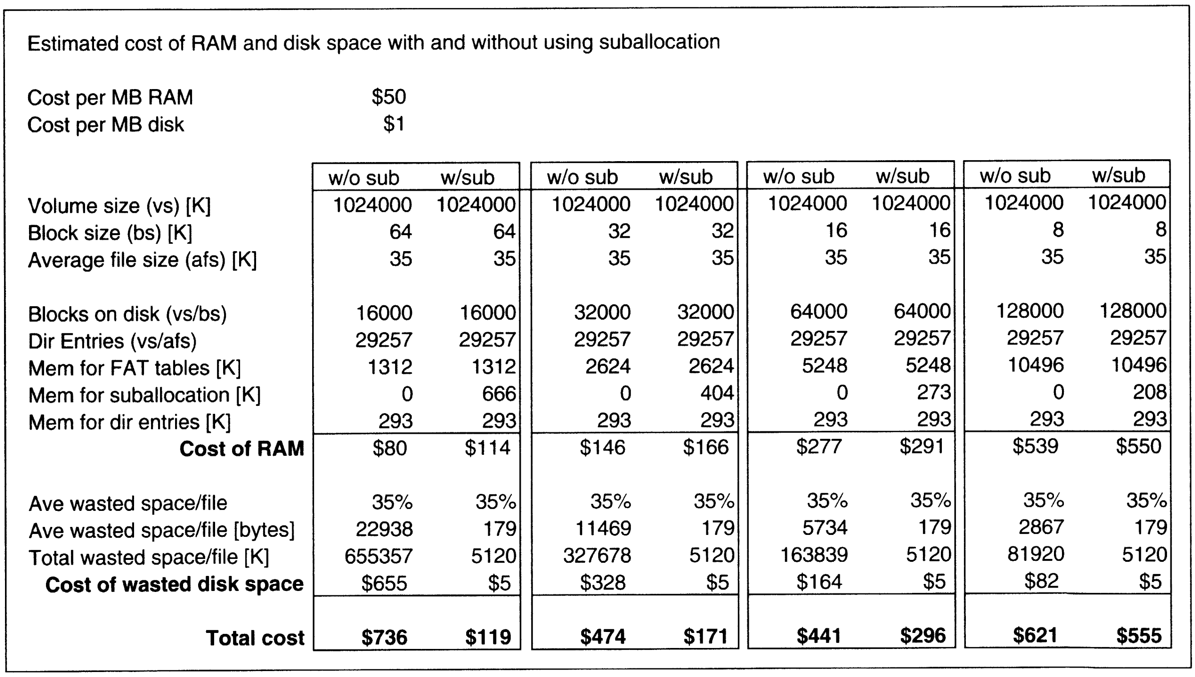

Estimated cost of RAM and disk space with and without using suballocation

Cost per MB RAM $50
Cost per MB disk $1

	w/o sub	w/sub	w/o sub	w/sub	w/o sub	w/sub	w/o sub	w/sub
Volume size (vs) [K]	1024000	1024000	1024000	1024000	1024000	1024000	1024000	1024000
Block size (bs) [K]	64	64	32	32	16	16	8	8
Average file size (afs) [K]	35	35	35	35	35	35	35	35
Blocks on disk (vs/bs)	16000	16000	32000	32000	64000	64000	128000	128000
Dir Entries (vs/afs)	29257	29257	29257	29257	29257	29257	29257	29257
Mem for FAT tables [K]	1312	1312	2624	2624	5248	5248	10496	10496
Mem for suballocation [K]	0	666	0	404	0	273	0	208
Mem for dir entries [K]	293	293	293	293	293	293	293	293
Cost of RAM	$80	$114	$146	$166	$277	$291	$539	$550
Ave wasted space/file	35%	35%	35%	35%	35%	35%	35%	35%
Ave wasted space/file [bytes]	22938	179	11469	179	5734	179	2867	179
Total wasted space/file [K]	655357	5120	327678	5120	163839	5120	81920	5120
Cost of wasted disk space	$655	$5	$328	$5	$164	$5	$82	$5
Total cost	**$736**	**$119**	**$474**	**$171**	**$441**	**$296**	**$621**	**$555**

Note: This spreadsheet and other server analysis utilities are available on NetWire, NOVLIB Library 11, in a self-extracting file named COMSUB.EXE. (See the Research Index for more information about downloading files from NetWire.)

Figure 10 charts the results for our 1GB disk, which had an average file size of 35KB. Not using suballocation cost us $736 in wasted storage and RAM. With suballocation enabled, the cost was only $119—a savings of $615. That's enough to buy a new 600MB hard disk!

Figure 10: Estimated costs of RAM and unused disk space with and without suballocation on a Novell production server.

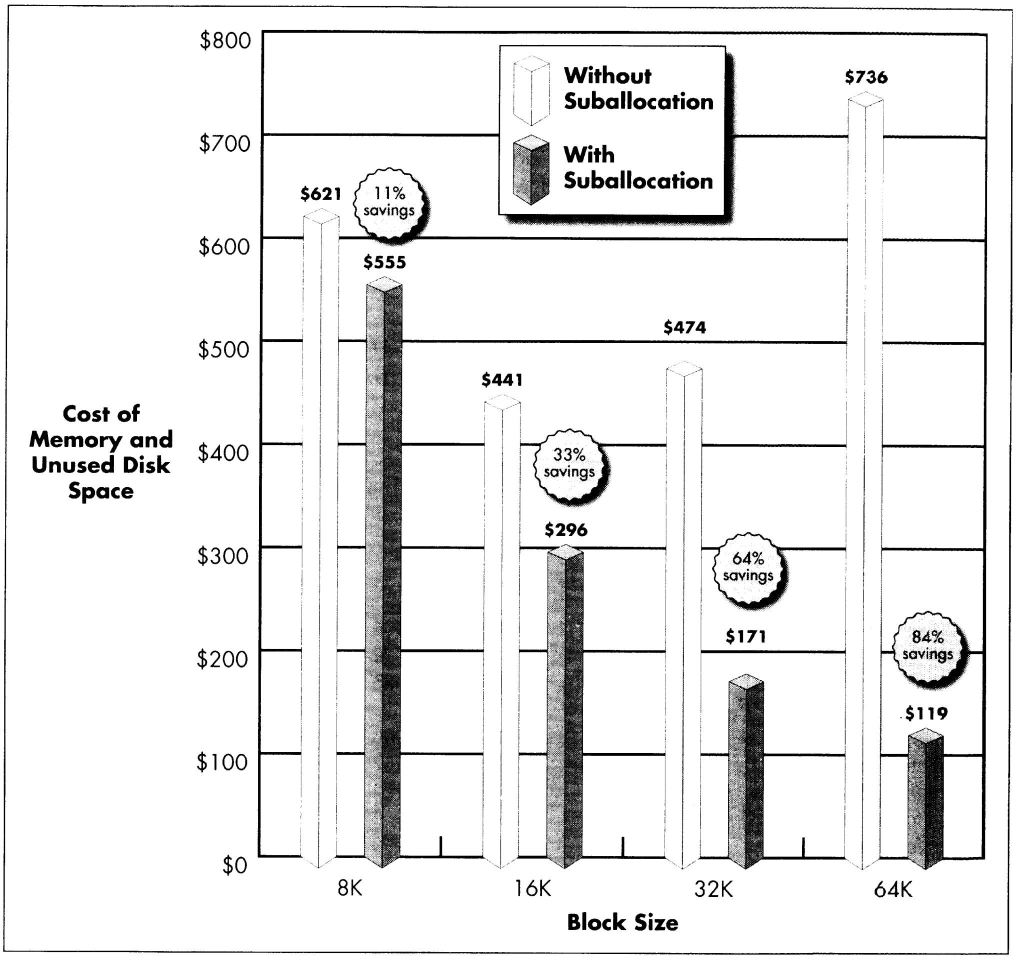

Enabling suballocation and increasing the disk block size results in lowering the overall file system memory requirements (especially for the File Allocation Tables) and increases throughput. Depending on the average file size and disk block size, suballocation can save you a lot of money on hardware. Put another way, suballocation gives you more storage for your buck (or franc or mark or pound).

Compression and Suballocation: Determining Your Savings

The most accurate way of predicting how much storage will be saved by implementing compression and suballocation is to analyze each file and compute (1) its compression size, and (2) its utilization within a disk block.

An easier method, which is just an estimate, involves the following steps:

1. Compute the amount of wasted space on the volume with NDIR, SDIR, or some other utility such as File Size from Norton Utilities (Symantec Corporation). To do this, subtract the total allocated space from the actual space used.

   Using `NDIR /C /UN` on a sample NetWare 3.12 volume with 32KB disk blocks, the results were "157,857,716 bytes in 4,180 files; 181,796,864 bytes in 5,548 blocks." So the wasted space was 181,796,864 – 157,857,716 = 23,939,148 bytes.

2. Now determine how much space will be saved by using suballocation and the existing block size. (If you will be increasing the block size, your savings will be even greater.) Using a worst-case scenario of 511 bytes of space wasted for each file, multiply the number of files by 511 and subtract from the wasted space (21,803,168 – 4,180 x 511 = 21,803,168).

3. Estimate the total space taken by executables (EXE, DLL) and multiply by .45 (assuming an average compression ratio of 45%). From the root of a volume, use the command `NDIR *.EXE,*.HLP /SUB /C /UN` and use the total.

4. Estimate the total space taken by all other files and multiply by .6 (assuming an average compression ratio of 60% for non-executable files).

5. Add the results from steps 2, 3, and 4 to estimate your total savings.

Summary

NetWare 4 provides new ways to effectively get more storage from your disks. By enabling compression and suballocation, data storage is better utilized. Infrequently-used and recently-deleted files no longer take up valuable disk real estate because they are compressed and suballocated. And those small batch files and documents no longer take up an entire disk block.

Perhaps the best feature of these technologies is that the user is shielded from the complexities of the system. And administrators can choose to accept the default server configuration and still make use of compression and suballocation.

APPENDIX C Answers to Exercises

This appendix contains the answers to the exercises in this course.

Exercise 1-1

Hands-on activity—no answers necessary.

Exercise 1-2

Hands-on activity—no answers necessary.

Exercise 1-3

Hands-on activity—no answers necessary.

Exercise 2-1

Group Activity

1. b
2. True
3. True
4. False
5. False
6. True
7. b
8. c
9. a
10. False
11. a
12. b
13. b
14. True
15. CX /R
16. True
17. b
18. a

Exercise 2-2

Scenario 1 suggested answer.

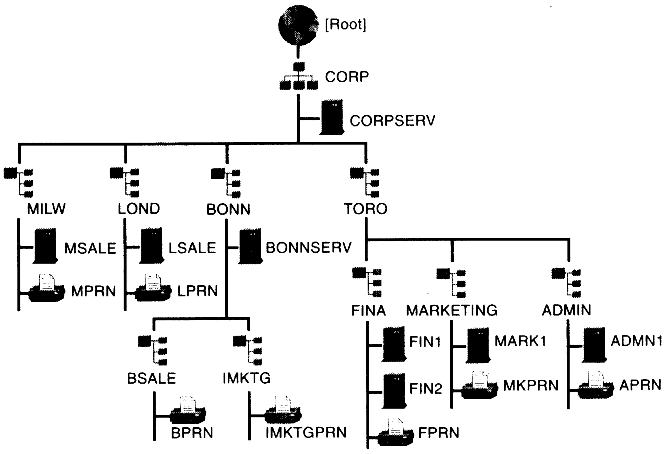

Scenario 2 suggested answer:

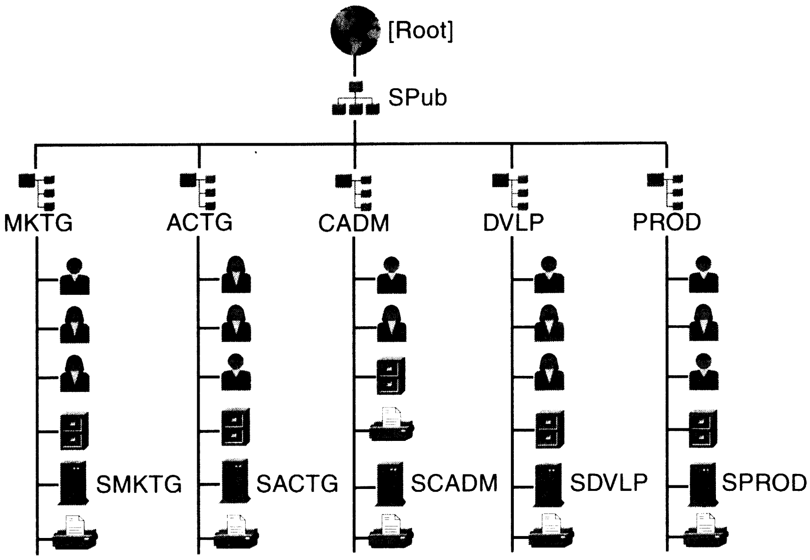

Scenario 3 suggested answer:

Before

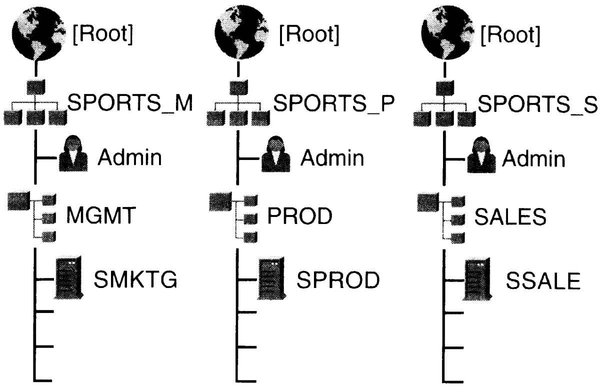

After

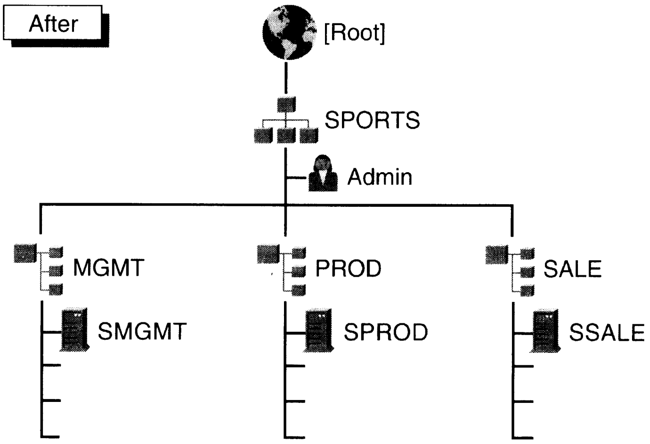

Exercise 2-3

1. Merged Directory tree

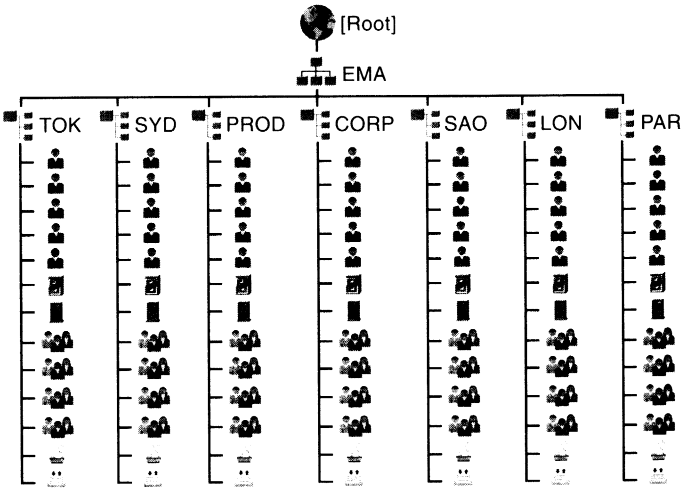

2. Suggested order:
   a. EMAPROD
   b. EMASAO
   c. EMALON
   d. EMAPAR
   e. EMASYD
   f. EMATOK

Exercise 2-4

Part I:

Answers will vary from class to class. Record the correct answers for your tree structure here.

1. [Root] and the Organization objects
2. No answer necessary
3. No answer necessary
4. Seven Organization objects
5. Organization
6. No answer necessary
7. Answer will vary
8. Answer will vary

Part II:

Hands-on activity—no answers necessary.

Exercise 2-5

1. .CN=LiPang.OU=CORP.O=EMA
2. .EPROD_P.PROD.EMA
3. .OU=CORP.O=EMA
4. KimH
5. .CN=EPROD_SYS.OU=PROD.
6. DougC
7. The current context is your workstation's current position in the Directory tree. The workstation can maintain only one current context per DOS session. You can set your current context to container objects only.
8. NAME CONTEXT = "PROD.EMA"

Exercise 3-1

Group Activity—no answers necessary.

Exercise 3-2

1. No answer required.
2. No answer required.
3. No answer required.
4. No answer required.
5. No answer required.
6. Supervisor object right gives implied Supervisor rights to All Properties property rights.
   a. Supervisor rights to All Properties.
   b. Supervisor object right.

      The Supervisor object right assigned to the Organizational role at the Organizational Unit object flows down to the Server object because the Server object is in the Organizational Unit.
   c. Yes. The Supervisor object right to the Server object gives the administrator Supervisor rights to the file system on all volumes attached to that Server object.
7. No answer required.
8. No answer required.
9. All rights including Supervisor.
10. No answer required.
11. No answer required.
12. An explicit rights assignment of Compare and Read property rights [CR] to the Object Trustees (ACL) property removes the Supervisor right to the Server object, which removes the Supervisor file system right from the volumes connected to the server.

    The explicit rights assignment also takes away the ability of the User object to create trustees for the Server object or assign Supervisor rights to himself or herself.

13. No answer required.

14. No answer required.

15. No answer required.

16. No answer required.

17. No answer required.

18. No answer required.

19. Browse object rights

    Yes. This will affect the user's file system rights. He or she will no longer have Supervisor rights in the file system.

Exercise 3-3

1. All rights, including Supervisor.

2. No answer.

3. No answer.

4. The Server object; no (MarcJ is not a trustee to the Server object).

5. No answer.

6. MarcJ receives file system rights from being a trustee of the container EMACORP and having Supervisor object rights.

7. No answer.

8. Groups, Organizational Roles, Security Equivalence, and rights given to the [Public] trustee and to the [Root].

9. RussC receives file system rights from having security equivalence to MarcJ, who has Supervisor object rights to the container EMACORP.

10. Place an IRF at the Server object ECORP to block the Supervisor right. This still gives MarcJ the ability to administrator the container but does not give him Supervisor file system rights.

11. Remove RussC's security equivalence to MarcJ.

Exercise 3-4

7. The auditor should change the password provided by the network administrator.

Exercise 4-1

1. No answer required.
2. No answer required.
3. No answer required.
4. No answer required.
5. The icon on the left means the container is the root of a partition. The icon on the right is the standard [Root] icon.
6. No answer required.
7. Replicas are stored on the following servers:
   - ECORP.EMA
   - ELON.LON.EMALON
   - EPAR.PAR.EMAPAR
   - EPROD.EMA
   - ESAO.SAO.EMASAO
   - ESYD.SYD.EMASYD
   - ETOK.TOK.EMATOK
8. Types of replicas:
   - ECORP.EMA = Master
   - ELON.LON.EMALON = Read/write
   - EPAR.PAR.EMAPAR = Read/write
   - EPROD.EMA = Read/write
   - ESAO.SAO.EMASAO = Read/write
   - ESYD.SYD.EMASYD = Read/write
   - ETOK.TOK.EMATOK = Read/write
9. No answer required.
10. No answer required.

11. No answer required.

12. The CORP Organizational Unit is not a partition root. Explain, if necessary, that the root-most container in a partition is called a partition root.

    EMASAO is a partition root. That is, the Directory has been divided so that EMASAO is the root-most container.

13. No answer required.

14. Only the Server object is displayed. In the Browser for Partition Manager, only objects that have partitioning capabilities are displayed. These are container objects and servers; containers can be partition roots, and servers can store replicas of partitions.

15. No answer required.

16. No answer required.

17. ECORP contains a master replica of the [Root] partition. It also contains a subordinate reference of the other partitions.

Exercise 4-2

1.

HR-ADMIN	Single Reference
HR-PERSONNEL	Secondary
HR-RECRUITING	Secondary

2. Leave it as the default, single reference server with Secondary servers.

3. In preparing for a merge.

Exercise 4-3

Hands-on activity—no answers necessary.

Exercise 5-1

Scenario 1 suggested solution:

Servers/Type	Partitions			
	[Root]	BONN	LOND	MILW
CORP1 Reference	Master	R/W	R/W	R/W
FINA1 Secondary	R/W			
MARK1 Secondary	R/W			R/W
LSALE Primary	R/W	R/W	M	
BSALE Primary		M	R/W	
MSALE Primary	R/W			M

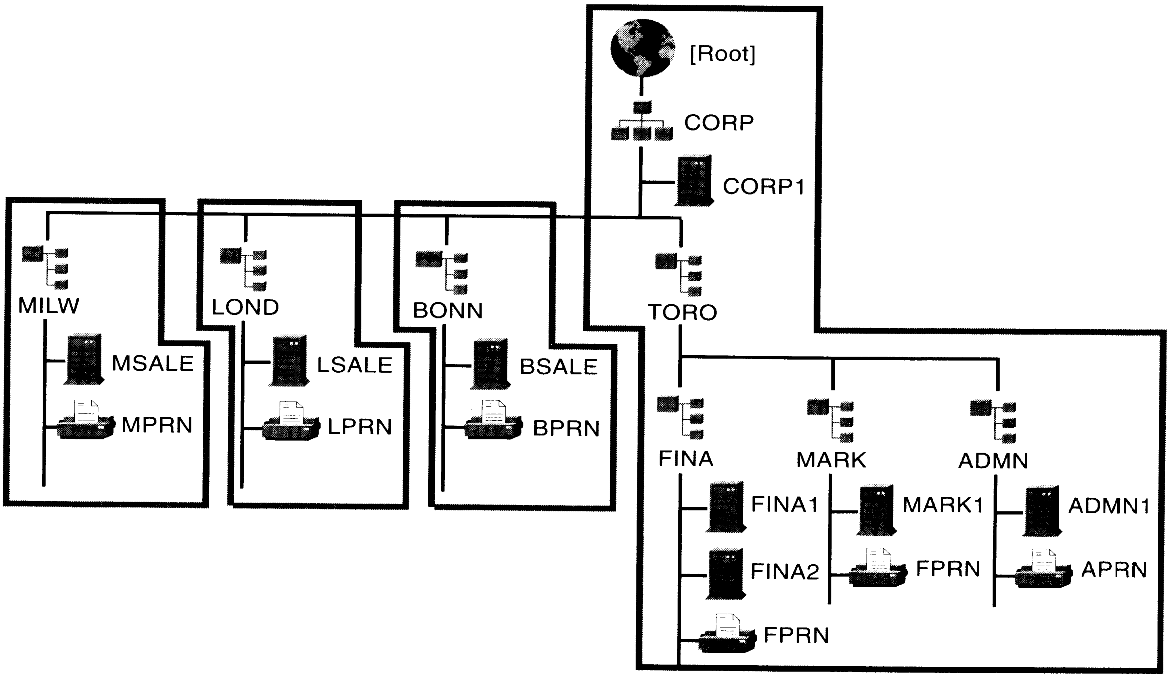

Scenario 2 suggested solution:

If the site for the company included several hundred servers and thousands of clients, you would make partitions by workgroups.

Servers	Time Type	Partitions		
		[Root]		
SMKTG	Single	Master		
SACTG	Secondary			
SCADM	Secondary	R/W		
SDVLP	Secondary			
SPROD	Secondary	R/W		

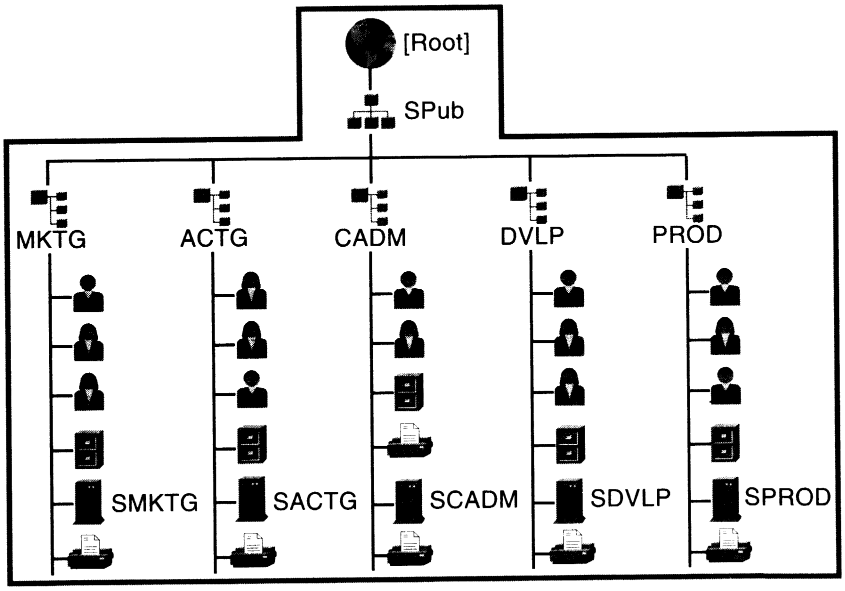

Scenario 3 suggested solution:

Servers	Time Type	Partitions		
		[Root]		
SPROD1	Single	Master		
SMGMT1	Secondary	R/W		
SSALE	Secondary	R/W		
SPROD2	Secondary			
SMGMT2	Secondary			

After

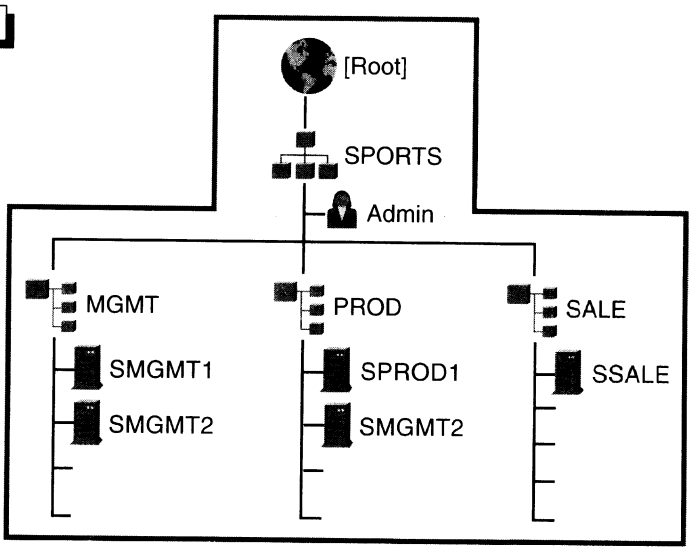

Exercise 6-1

1. No answer necessary
2. No answer necessary
3. No answer necessary
4. No answer necessary
5. No answer necessary
6. No answer necessary
7. No answer necessary
8. No answer necessary
9. NLIST PRINTER

   Printer objects do not exist in a NetWare 3.1*x* bindery.
10. No answer necessary
11. NETADMIN would not load because students are connected to the server in bindery mode, not NDS mode.
12. No answer necessary
13. No answer necessary
14. You can view and manage only the bindery-related information.
15. You can view and manage only the bindery-related information.
16. No answer necessary
17. No answer necessary
18. No answer necessary
19. No answer necessary

Exercise 8-1

1. LOAD INSTALL at the server console.
2. Load the INSTALL.NLM.

   Select Product options.
3. Postmaster General
4. Full distinguished name (Example: Admin.CORP)
5. Messaging Server object

   Message Routing Group object
6. SYS:\PUBLIC
7. \MHS\MAIL\USERS\username\
8. LOAD MHS
9. AUTOEXEC.NCF

Exercise 8-2

Hands-on activity—no answers necessary.

Exercise 8-3

Hands-on activity—no answers necessary.

Exercise 9-1

Hands-on activity—no answers necessary.

Exercise 9-2

Answers for MONITOR Information:

1. Utilization is 85 percent.
2. Determine if the network board is performing adequately or if an ill-behaved NLM exists.

Answers for Cache Utilization:

1. Long Term Cache Hits is low.
2. Add more RAM, unload noncritical NLMs, or use the REMOVE DOS command to free memory.

Answers for Custom Statistics:

1. Three parameters show a value of 100.
2. Upgrade the network board.

Answers for Polling/Processor Utilization:

1. The polling percentage is high.
2. No.
3. High.
4. Add more RAM, unload noncritical NLMs, move some applications to another server, or upgrade the CPU.

Answers for Memory Information:

1. Percent Free is very low.
2. It could be caused by ill-behaved NLMs; isolate the problem NLMs.

   Too many modules are loaded; unload some NLMs or add more RAM to the server.

Exercise 9-3

1. Increased throughput; therefore, increased performance.
2. Probably not. This parameter mainly affects larger files.
3. Yes. Increased throughput may be significant because of the reduction of the large number of reads and writes occurring with small files.
4. Probably not. LIP would only have an effect with larger files.

APPENDIX D Study Material for Exercise 2-1

Introduction

Your instructor will assign you to a group.

You will have 10 minutes to study the section listed below that corresponds to your group.

You will be required to teach others what you learn.

- Group 1: Study pages D-2 through D-4.
- Group 2: Study pages D-5 through D-6.
- Group 3: Study pages D-7 through D-9.
- Group 4: Study pages D-10 through D-12.

Group 1

Planning a Directory Tree

Planning an efficient Directory tree can

- Make looking up information easier for users
- Make administering the network easier for network supervisors
- Provide fault-tolerance for the Directory database
- Decrease network traffic

In general, the Directory tree should be organized based on a logical, and not a physical, organization of shared network resources. This means that objects representing shared network resources should be placed in common containers. The structure of the Directory tree can also be based on your organizational structure, geographic location, or administrative responsibilities, or a combination of all of these.

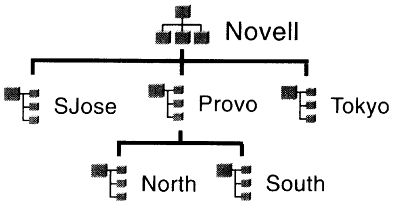

Figure D-1: Directory Tree Based on Geographic Location

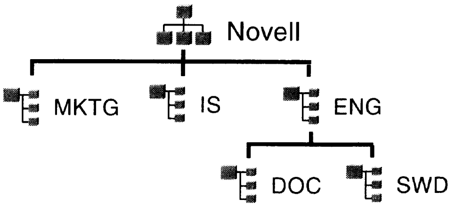

Figure D-2: Directory Tree Based on Organizational Structure

Many factors can influence the structure of your Directory tree. You may need to study workgroups, resource allocation, and information flows within your organization to determine how to organize the Directory objects.

Extensive planning is not required. NDS is easily reconfigured and modified.

How the Directory Tree Affects Resource Access

A multiple-container Directory tree probably has the greatest impact on how network resources are accessed because of the following:

- Objects are no longer in a single container.
- NDS does not search the entire Directory tree to find an object.

For these reasons, NDS requires precise information to find the correct object. In Figure D-3, for example, three User objects named Bob exist in separate containers in the Directory. If you entered LOGIN BOB, which User object would NDS use?

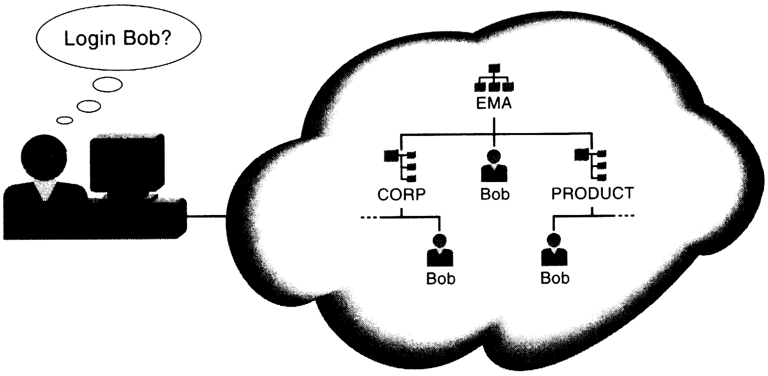

Figure D-3: Which Object Gets Used?

You, or your workstation, must provide NDS with enough information to locate the object in the Directory tree. This information comes in the form of an appropriate object name.

Said another way, to correctly access network resources, you must use the correct object name; the correct object name exactly identifies which object in the entire Directory tree you want.

This information can be provided by using either of the following:

- Distinguished name
- Relative distinguished name

Object-Naming Terminology

To understand the difference between distinguished name and relative distinguished name, you must be familiar with the terminology associated with naming objects. The following subsections explain NDS object-naming terminology.

Common Name

A leaf object's *common name* (CN) is the name shown next to the leaf object in the Directory tree.

A common name is a relative distinguished name.

For example, the common name for both User objects in Figure D-4 is MarcJ.

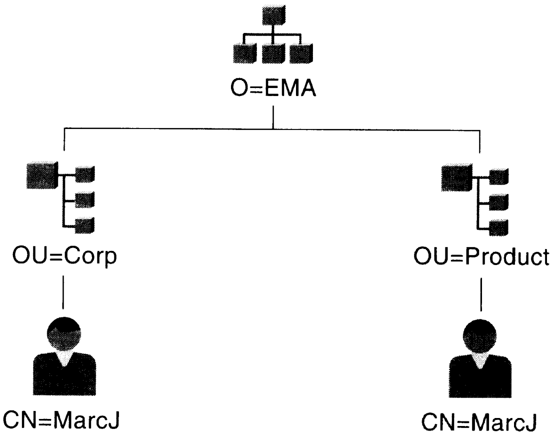

Figure D-4: Common Names (CNs) for Leaf Objects

Context

Context is an object's position in the Directory tree. It is a list of the container objects leading from the object to the [Root]. (Locating an object through context is similar to locating a file using the directory path.)

In Figure D-4 (above), the difference between the two MarcJ User objects is their context. The User object on the left is in OU=Corp.O=EMA; the User object on the right is in OU=Product.O=EMA.

Group 2

Distinguished Name

An object's *distinguished name* is a combination of its common name and its context.

For example, in Figure D-5, the distinguished name for the User object MarcJ in the Organizational Unit Corp in the Organization EMA is as follows:

.CN=MarcJ.OU=Corp.O=EMA

The distinguished name for the User object MarcJ in the Organizational Unit Product in the Organization EMA is as follows:

.CN=MarcJ.OU=Product.O=EMA

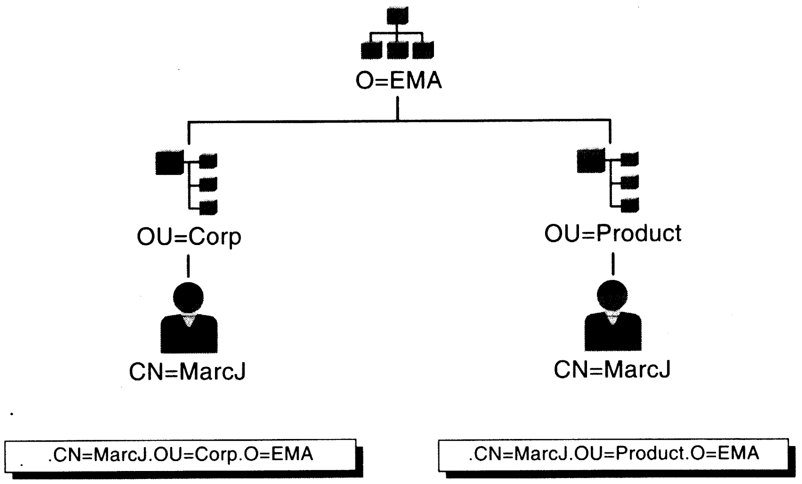

Figure D-5: Distinguished Name Examples

A distinguished name starts with a leading period. The objects in the name are separated by periods, similar to the backslash used in DOS paths. Trailing periods cannot be used.

An object is exactly identified with a distinguished name. Two objects cannot have the same distinguished name.

Current Context

Your current context can affect how much of an object's distinguished name you must provide with a command to access the resource.

Current context (sometimes called name context) can be thought of as your current position in the Directory tree. It is a pointer in the NetWare DOS Requester™, similar to a network drive, that identifies the default NDS container for your workstation.

You can refer to an object in your current context by its common name since your current context and the object's context are the same.

Relative Distinguished Name

A *relative distinguished name* lists the path of objects leading from the object to the current context.

A relative distinguished name does *not* start with a leading period, but the objects in the name are separated by periods. Trailing periods (explained on the next page) can also be used.

In Figure D-6, for example, you could refer to each MarcJ User object as follows if your current context were O=EMA:

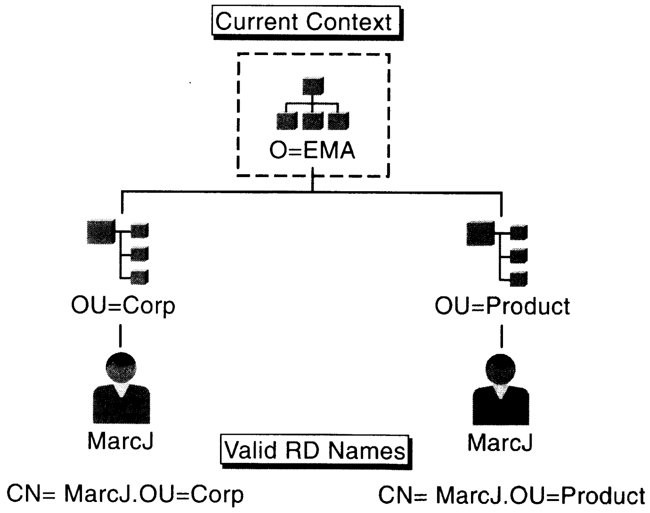

Figure D-6: Relative Distinguished Name Examples

When you use a relative distinguished name, NDS must build a distinguished name from it. This is accomplished by appending the relative distinguished name to the current context.

> Relative Distinguished Name + Current Context = Distinguished Name

In the following examples, a different current context creates a different distinguished name when the same relative distinguished name is submitted:

[illegible]	Current Context	Resulting Distinguished Name
CN=MarcJ	O=EMA	.CN=MarcJ.O=EMA
CN=MarcJ	OU=Corp.O=EMA	.CN=MarcJ.OU=Corp.O=EMA

Table D-1: How Relative Distinguished Names Work

Group 3

Trailing Periods

You can only use *trailing periods* in relative distinguished names. Each *trailing period* tells NDS to remove one object name from the left side of the current context.

For example, if your current context is

OU=Acctg.OU=Corp.O=EMA

and you enter the relative distinguished name

CN=ADMIN..

NDS will remove two names from the left side of the current context and append the relative distinguished name to the remaining name. This action would produce the following distinguished name:

.CN=Admin.O=EMA

Typeful Naming

A *typeful name* uses attribute type abbreviations (see Table D-2 below) to distinguish between the different container types and leaf objects in an object's distinguished or relative distinguished name. While not mandatory, attribute types help to avoid the confusion that can occur with typeless naming (explained on the next page).

Table D-2 lists the attribute type abbreviations. Note that all leaf objects use the *Common Name* attribute type abbreviation.

Attribute Type Abbreviation	Object
C	Country
O	Organization
OU	Organizational Unit
CN (Common Name)	All leaf objects

Table D-2: Attribute Type Abbreviations

For example, the following is a typeful name:

.CN=MarcJ.OU=Product.O=EMA

Typeless Naming

A *typeless name* does not include the object attribute type. For example, the typeless distinguished name for .CN=MarcJ.OU=Product.O=EMA is as follows:

.MarcJ.Product.EMA

If you do not provide a typeful object name, NDS will calculate attribute types for each object.

Accessing an Object by Its Common Name in Utilities

Objects in your current context can be referred to by their common names.

For example, consider the EMA Directory tree shown below. If your current context is OU=Product.O=EMA, any of the following commands would work:

```
LOGIN DOUGC
MAP S1:=DBAPP
CAPTURE P=PROD_P
```

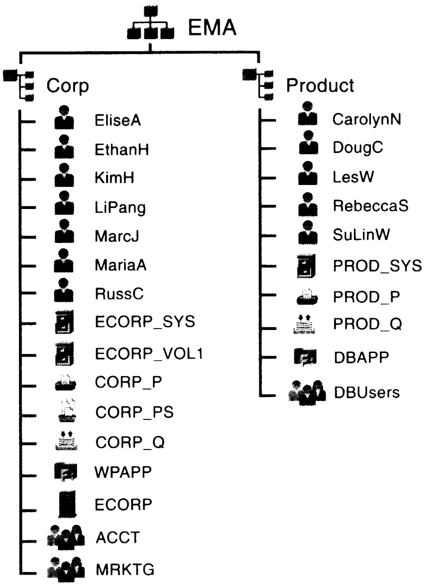

Figure D-7: The EMA Directory Tree

If your current context is OU=Corp.O=EMA and you want to access an object in Product, you must first make Product your current context or use a distinguished name for the object.

For this reason, a user's current context should be set to the container that holds the resources he or she most commonly uses.

Changing Your Current Context with CX

Accessing an object by its common name requires that you set or change your current context to the object's context beforehand.

CX allows you to determine your current context, view containers in a Directory tree structure, and change your context to a specific container object.

The following table shows how to perform several tasks with CX.

To	Enter
View your current context	**CX**
View the Directory tree structure below your current context	**CX /T**
View all objects in your current context	**CX /All /T**
Change your current context using a distinguished name	**CX .OU=Product.O=EMA**
Change your current context to the [Root]	**CX /R**
Change your current context to your parent container	**CX .**

Table D-3: Common Tasks Accomplished with CX

Group 4

Accessing Objects Using Distinguished Names

A typeful distinguished name exactly identifies an object and contains all the information NDS needs to locate the object in the Directory tree.

For example, consider the EMA Directory tree shown below. If you want to access the objects in .OU=Product.O=EMA, any of the following commands would work, regardless of your current context:

```
LOGIN .CN=DOUGC.OU=PRODUCT.O=EMA
MAP S1:=.CN=DBAPP.OU=PRODUCT.O=EMA
CAPTURE P=.CN=PROD_P.OU=PRODUCT.O=EMA
```

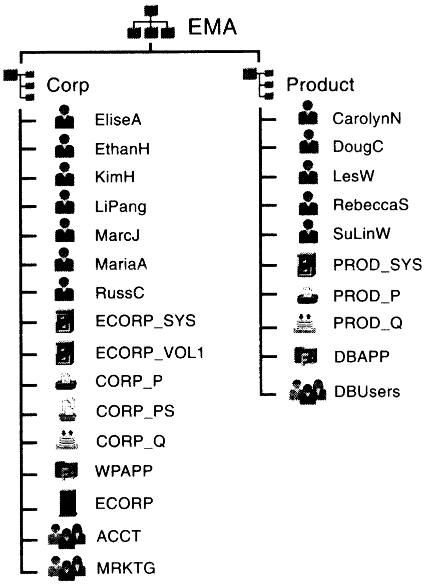

Figure D-8: The EMA Directory Tree

Setting the User's Current Context at Login

The user's current context can be set before or during login.

Setting the Current Context Before Login

You can set a workstation's current context before login by adding this statement to the NET.CFG file:

NAME CONTEXT=*"distinguished name path"*

For example:

NAME CONTEXT="OU=Marketing.O=EMA"

When the NetWare DOS Requester is loaded (by executing VLM.EXE), it reads the NET.CFG file and sets your workstation's current context to the context in the NAME CONTEXT statement.

This method of setting current context is workstation specific and cannot be set differently for individual users on the same workstation.

Setting the Current Context During Login

You can set a workstation's current context during login by adding this command to a container, profile, or user login script:

CONTEXT *distinguishedname*

For example:

CONTEXT .PROD.EMA

This method of setting current context is not workstation specific; it can be set for an individual or a group.

Using Correct Naming in Login Scripts

The login process is affected by a workstation's current context if objects used in the login script use common or relative distinguished names.

When creating login scripts in a multiple container environment, you need to identify each resource appropriately:

- Use distinguished names to access the same object at each login, regardless of the current context at login.
- Use common names or other relative distinguished names to access objects relative to the current context of the workstation at login.

For example, the following login script commands will always access resources in Product.EMA, regardless of the current context at login:

```
MAP INS S1:=.DPAPP.PRODUCT.EMA

CAPTURE P=.PROD_P.PRODUCT.EMA
```

In contrast, the following commands will access resources in either Corp.EMA or Product.EMA, depending on the workstation's current context:

```
MAP INS S1:=WPAPP

MAP INS S2:=SSAPP
```

Notes

Glossary

The following glossary entries are based on the *NetWare 4.1 Concepts* manual. For more complete information about a particular term, refer to the same term in the *Concepts* manual.

Abend

(Abnormal end) A message issued by the operating system when it detects a serious problem, such as a hardware or software failure. The abend stops the NetWare® server.

Abend messages are explained in *System Messages*.

Accelerated

A file open mode that provides a fast response time when an application updates data files. However, using Accelerated mode to open a file created with Btrieve 5.*x* or earlier disables Btrieve's automatic data recovery capability. *See also* Exclusive, file open mode, Normal, Read-Only, and Verify.

Access Control List

(ACL) An object property that stores information about who or what can access that object. An ACL contains trustee assignments that include object and property rights. The ACL also contains the Inherited Rights Filter. When you view an object's trustees or its Inherited Rights Filter, you are seeing the *values* of that object's ACL.

An ACL for an object is like the list of trustees for a file or directory. To change an ACL (and therefore a trustee's rights to an object), you must have a property right that allows you to modify the ACL for that object.

See also Object; Security.

Access Control right

A file system right that grants the right to change the trustee assignments and Inherited Rights Filter of a directory or file.

See also Rights.

access path

See index path.

Accounting

The process of tracking resources used on a network. The network supervisor can charge for network services and resources by assigning users account balances that they draw from as they use the services and resources.

ACL

(Access Control List) An object property that stores information about who or what can access that object.

See also Access Control List.

Add or Delete Self right

A property right that grants a trustee the right to add or remove itself as a value of the property.

See also Rights.

Add-on board

A circuit board that modifies or enhances a personal computer's capabilities. Examples include

- **Memory board.** Increases the amount of RAM within a personal computer.
- **Network board.** Allows workstations to communicate with each other and with the NetWare server. (Network boards are sometimes referred to as NIC cards.) Network boards are connected together with cabling.

Address

A number that identifies a location in memory or disk storage or that identifies the location of a device on the network.

See Controller address; IP address; Network numbering; SCSI.

Address Resolution Protocol

(ARP) A process in Internet Protocol (IP) and AppleTalk* networks that allows a host to find the Media Access Control (MAC) address of a target host on the same physical network when it only knows the target's IP address.

Addressing (disk channel)

The method of assigning numbers to identify hardware resources on disk channels. Each controller must have a unique address on the disk channel. You can find the physical address settings in the documentation shipped with the controller.

Addressing space

The total amount of RAM available to the operating system in a NetWare 4™ server. This amount can be divided into domains.

In the NetWare 4 operating system, the maximum addressing space is 4 gigabytes (GB), although practical hardware limits are much lower.

See also Paging.

ADSP

See AppleTalk Data Stream Protocol.

AFP

An AppleTalk protocol that provides communication and data transmission between file servers and clients in an AppleShare network.

See also AppleTalk Filing Protocol.

AFP Server object

A leaf object that represents an AppleTalk Filing Protocol server that is operating as a node on your NetWare network. The AFP server is probably also acting as a NetWare router and an AppleTalk server for several Apple* Macintosh* computers.

See also Object.

aggregate field value

A field value that a group aggregate function determines. This value is based on a set of field values selected from a table. *See also* group aggregate function.

aggregate function

See group aggregate function.

Alias object

A leaf object that points to the original location of an object in the Directory. Aliases can make NetWare Directory Services™ (NDS) easier to use. By using an alias, you can make any Directory object located in one place in the Directory also appear to be in another place in the Directory.

See also Object.

AppleShare software

Networking software, from Apple Computer, Inc., that enables a Macintosh computer to function as a file server in an AppleTalk network. Also, AppleShare* workstation software that allows a Macintosh computer to access an AppleShare server.

AppleTalk Filing Protocol

(AFP) An AppleTalk protocol that provides communication and data transmission between file servers and clients in an AppleShare network. When AFP.NLM is loaded on a NetWare server running NetWare for Macintosh, AFP allows Macintosh users to share files by interacting directly with the NetWare file system on the same level as NetWare Core Protocol™ (NCP).

AppleTalk Phase 2

The latest version of the AppleTalk protocols. AppleTalk Phase 2 implements more efficient routing techniques that improve performance in multiprotocol environments.

AppleTalk Print Services module

A NetWare Loadable Module™ (NLM) program that enables Macintosh users to print to NetWare queues and enables non-Macintosh users to print to AppleTalk printers.

AppleTalk protocols

The underlying forms and rules that determine communication between nodes on an AppleTalk network. These protocols include

- Link Access Protocols (LAPs)
- Datagram Delivery Protocol (DDP)
- Routing Table Maintenance Protocol (RTMP)
- AppleTalk Update-Based Routing Protocol (AURP)
- Name Binding Protocol (NBP)
- Printer Access Protocol (PAP)
- Zone Information Protocol (ZIP).

These protocols control the AppleTalk network, from the network board to the application software.

See also AppleShare software; AppleTalk Filing Protocol; AppleTalk Phase 2; Zones.

Application

A software program that makes calls to the operating system and manipulates data files, allowing a user to perform a specific job (such as accounting or word processing).

- **Standalone application.** An application that runs from the hard disk or floppy disk in a self-contained, independent computer. Only one user can access the application.
- **Network application.** An application that runs on networked computers and can be shared by users.

  Network applications use network resources such as printers. Advanced network applications (such as electronic mail) allow communication among network users.

Archive

A transfer of files to long-term storage media, such as optical discs or magnetic tape.

See also Attributes; Backup; Data migration; High Capacity Storage System; Storage Management Services.

Archive Needed (A) attribute

A file attribute, set by the NetWare operating system, indicating that the file has been changed since the last time it was backed up.

See also Attributes.

ARP

(Address Resolution Protocol) A process in IP and AppleTalk networks that allows a host to find the MAC™ address of a target host on the same physical network when it only knows the target's IP address.

See also Address Resolution Protocol.

ascending

The default collating order for an index. *See also* descending.

ATPS

(AppleTalk Print Services) Enables a Macintosh client to print to a NetWare queue; also enables a non-Macintosh client to print to an AppleTalk printer.

See also AppleTalk Print Services module.

Attach

Establishes a connection between a workstation and a NetWare server.

In NetWare 4, with NetWare Directory Services, users no longer need to attach separately to multiple servers. When users log in to the Directory tree, they automatically have access to any resources in the Directory tree to which they have rights. Rights to resources are verified through authentication. The ATTACH command can still be used in login scripts to establish connections with bindery-based servers.

See also Authentication; NetWare Directory Services.

Attributes

The characteristics of a directory or file. In NetWare, these characteristics are called flags. Attributes dictate what can be done with a file or directory. For example, you can set a file to be a Read Only (Ro) file. You can set or clear attributes with the FLAG command line utility, the FILER menu utility, or the NetWare Administrator graphical utility.

Note: NDS objects do not have attributes.

Auditing

The process of examining network transactions to ensure that network records are accurate and secure. NetWare auditing allows individuals, acting independently of network supervisors and other users, to audit network transactions.

Authentication

A means of verifying that an object sending messages or requests to NDS is authorized to do so.

AUTOEXEC.BAT

A batch file that executes automatically when DOS or OS/2 is booted on a computer. A workstation's AUTOEXEC.BAT file, located on the bootable floppy or hard disk, can contain commands that

- Load NetWare client files
- Load other files required by the hardware
- Set the DOS or OS/2 prompt
- Change the default drive to the first network drive
- Log in the user

The workstation AUTOEXEC.BAT file can also load user-specific programs such as NETBIOS.COM or call other batch files.

A NetWare server's AUTOEXEC.BAT file, located on the hard disk's DOS or OS/2 partition, can contain the command that loads the NetWare operating system (SERVER.EXE).

AUTOEXEC.NCF

A NetWare server executable batch file, located on the NetWare partition of the server's hard disk. AUTOEXEC.NCF is used to

- Load modules
- Set the NetWare operating system configuration
- Set bindery contexts
- Store the IPX™ internal network number
- Store the file server name
- Make time zone settings

The network supervisor can also add executable server commands (such as LOAD INSTALL or LOAD MONITOR) to AUTOEXEC.NCF.

Automatic rollback

A feature of TTS that returns a database to its original state. When a network running under TTS fails during a transaction, the database returns, or rolls back, to its most recent complete state, preventing corruption from an incomplete transaction.

See also Transaction Tracking System.

Backup

A duplicate of data (file, directory, volume), copied to a storage device (floppy diskette, cartridge tape, hard disk). A backup can be retrieved and restored if the original is corrupted or destroyed.

See also Data set.

Backup hosts and target

A backup host is a NetWare server that has a storage device and a storage device controller attached.

A target is the server, workstation, or database from which you back up data or to which you restore data.

See also Backup; Storage Management Services; Target Service Agent.

Baud rate

In serial communication, the signal modulation rate, or the speed at which a signal changes.

See also Serial communication.

Bindery

A network database, in NetWare versions earlier than NetWare 4, that contains definitions for entities such as users, groups, and workgroups. In NetWare 4, the bindery has been replaced by the Directory database. Bindery services provides NetWare 4 networks with backward compatibility to NetWare versions that used the bindery.

See also Bindery services.

Bindery context

The container object in which bindery services is set.

See also Bindery services.

Bindery context path

A path statement that allows bindery context to be set in as many as 16 containers. Use the SET parameter to set bindery contexts. Multiple contexts are separated by semicolons. For example:

```
SET BINDERY CONTEXT = OU=Legal.O=Novell;
OU=Sales.O=Novell; OU=Mktg.O=Novell
```

See also Bindery services.

Bindery object

A leaf object that represents an object placed in the Directory tree by an upgrade or migration utility; NDS cannot identify the object. This object provides backward compatibility with bindery-oriented utilities.

See also Object.

Bindery Queue object

A leaf object that represents a queue placed in the Directory tree by an upgrade or migration utility; NDS cannot identify the object. This object provides backward compatibility with bindery-oriented utilities.

See also Object.

Bindery services

A feature of NetWare 4 that allows bindery utilities and clients to co-exist with NDS on the network. Objects in a bindery exist in a flat database instead of a hierarchical database (such as a Directory tree). Bindery services creates a flat structure for the objects within an Organization or Organizational Unit object.

See also Context; Directory tree; Object.

Binding and unbinding

The process of assigning (binding) a communication protocol to network boards and LAN drivers and the process of removing (unbinding) it. Each network board must have at least one communication protocol bound to the LAN driver for that board. Without a communication protocol, the LAN driver cannot process packets.

You can bind more than one protocol to the same LAN driver and board. You can also bind the same protocol stack to multiple LAN drivers on the server. You can also cable workstations with different protocols to the same cabling scheme.

See also Communication protocols; IPX external network number.

BIOS

(Basic Input/Output System) A set of programs, usually in firmware, that enable each computer's central processing unit to communicate with printers, disks, keyboards, consoles, and other attached input and output devices.

Block

The smallest amount of disk space that can be allocated at one time on a NetWare volume. The block size depends on the size of the volume. The block size is set automatically during installation; we recommend that you use the default block size. Block suballocation can subdivide a disk block among several files to make better use of disk space when a large block size is used.

See also Block suballocation.

Block suballocation

Allows parts of several files to share one disk block, better utilizing disk space. Block suballocation divides any partially used disk block into 512-byte suballocation blocks. These suballocation blocks are used to share the remainder on the block with leftover fragments of other files. Block suballocation is set by default when NetWare 4 is installed.

See also Block; Volume.

Boot files

Files, such as AUTOEXEC.BAT and CONFIG.SYS, that

- Start the operating system and its drivers
- Set environment variables
- Load NetWare

NetWare server boot files include AUTOEXEC.NCF and STARTUP.NCF. Workstation boot files depend on the client type (DOS, MS Windows, OS/2, Macintosh, or UNIX®).

BOOTP

A protocol used by some hosts to obtain their IP addresses.

Bridge

A device that retransmits packets from one segment of the network to another segment.

A router, on the other hand, is a device that receives instructions for forwarding packets between topologies and determines the most efficient path.

See also Router.

Browse right

An object right that grants the right to see an object in the Directory tree.

See also Rights.

Browsing

A way of finding objects in the Directory. Objects in the Directory are in hierarchical order. Since the Directory can be very large, you can browse the tree structure to find the object you need.

Btrieve

A complete key-indexed record management system designed for high-performance data handling.

Buffer

An area in server or workstation memory set aside to temporarily hold data, such as packets received from the network.

See Cache buffer; Packet receive buffer.

Cabling system

Part of a network's physical layout.

See also Topology.

Cache buffer

A block of NetWare server memory (RAM) in which files are temporarily stored. Cache buffers greatly increase NetWare server performance. The cache buffer size depends on the default block size, which depends on the size of the volume. (*See also* Block.) Cache buffers allow workstations to access data quickly; reading from and writing to memory is much faster than reading from or writing to disk.

See also Cache buffer pool.

Cache buffer pool

The amount of memory available for use by the operating system after the SERVER.EXE file has been loaded into memory. When an NLM is removed from server memory, the NLM returns the borrowed memory to the cache buffer pool.

Cache memory

Available RAM that NetWare uses to improve NetWare server access time. Cache memory allocates memory for the hash table, the FAT, the Turbo FAT, suballocation tables, the directory cache, a temporary data storage area for files and NLM files, and available memory for other functions. If the cache memory uses the default block size and a file takes more than one block, the file is placed in a second noncontiguous block both in cache memory and on the volume (on the hard disks).

See also Block; Directory hashing.

Can't Compress (Cc) attribute

A status flag indicating that a file cannot be compressed because of insignificant space savings.

See also Attributes.

CDM

(Custom Device Module) The driver component in the NetWare Peripheral Architecture™ (NPA) used to drive specific storage devices attached to the host adapter.

Channel

The path that data flows on to get from the computer to the device. This path can include a host bus adapter, cables, and storage devices.

Character length

In serial communication, the number of bits used to form a character.

See also Serial communication.

Child VLM

A Virtual Loadable Module™ (VLM) that handles a particular implementation of a logical grouping of functionality. For example, each of the NetWare server types has its own child VLM™:

- BIND.VLM for NetWare 2 and 3 bindery servers
- NDS.VLM for NetWare 4 NDS servers
- PNW.VLM for Personal NetWare™ (NetWare desktop) servers

Also, various implementations of transport protocols each have an individual child VLM. For example, IPXNCP.VLM handles IPX protocol services.

Client

A workstation that uses NetWare software to gain access to the network. In NetWare, client types include DOS, Macintosh, OS/2, UNIX, and MS Windows. With the respective client software, users can perform networking tasks. These tasks include mapping drives, capturing printer ports, sending messages, accessing files, and changing contexts.

See also DOS client; Macintosh client; OS/2 client; Workstation.

CMOS RAM

Memory used to store system configuration data (such as number of drives, types of drives, and amount of memory). CMOS RAM is battery powered to retain the date, time, and other information that needs to be maintained when the computer is turned off.

Code page

A table storing a character set that supports one or more language scripts. Many personal computers use operating systems that support multiple code pages and allow you to switch between them.

COM ports

Asynchronous serial ports on IBM PC-compatible computers.

See also Serial port.

Command format

Instructions that show how to type a command at the keyboard; also called *syntax*. In NetWare manuals, a command format may include constants, variables, and symbols.

Communication buffer

(Formerly used to refer to a *packet receive buffer*) An area in the NetWare server's memory set aside to temporarily hold data packets arriving from workstations.

See also Packet receive buffer.

Communication protocols

Conventions or rules used by a program or operating system to communicate between two or more endpoints. All communication protocols allow information to be packaged, sent from a source, and delivered to a destination system.

See also Binding and unbinding; Open Data-Link Interface.

Compare right

A property right that grants the right to compare any other value to a value of that property.

See also Rights.

Compressed (Co) attribute

A status flag indicating that a file is compressed.

See also Attributes.

Computer object

A leaf object that represents a computer on the network. The Computer object's properties can store information such as the computer's serial number or the person the computer is assigned to.

See also Object.

Configuration (hardware)

The equipment used on a network (such as servers, workstations, printers, cables, network boards, and routers) and the way the equipment is connected—the network's physical layout.

Configuration (router)

The settings and parameters chosen through internetwork utilities to configure a NetWare 4 server as a router.

Configuration (server)

The settings and parameters chosen through INSTALL.NLM while either installing a new NetWare 4 server or performing maintenance work on an existing NetWare 4 server.

Configuration (software)

The software used on a network for servers, clients, protocols, services, drivers, utilities, etc.; provides the means to communicate and operate on network hardware.

Connection number

A number assigned to any workstation that attaches to a NetWare server; it may be a different number each time a workstation attaches. Connection numbers are also assigned to processes, print servers, and applications that use server connections. The server's operating system uses connection numbers to control each workstation's communication with other workstations.

Connectivity

The ability to link computer systems (Macintoshes, PCs, minicomputers, and mainframes) into a network in order to share resources such as applications and printers.

See also Internetwork.

Console

The monitor and keyboard where you view and control NetWare server activity.

See also Server console.

Console operator

A user or member of a group who has been delegated rights to manage the NetWare server.

Container login script

Sets a general environment for all users in a container (such as an Organizational Unit). These login scripts execute first.

See also Login scripts.

Container object

An object that can hold, or contain, other objects. For example, an Organizational Unit is a container object because it can contain other objects. Container objects are used to logically organize all other objects in the Directory tree.

See also Object.

Context

The position of an object within its container in the Directory tree. NDS allows you to refer to objects according to their positions within a Directory tree. When you add an object (such as a server or user) to the network, you place that object in a container object in the Directory tree.

See also Container object; NetWare Directory Services; Object.

Controller address

The number the operating system uses to locate the controller on a disk channel. The number is physically set (usually with jumpers) on a disk controller board.

Controller board

A device that enables a computer to communicate with another device, such as a hard disk or tape drive. The controller board manages input/output and regulates the operation of its associated device.

See also SCSI.

Copy Inhibit (Ci) attribute

A file attribute, valid only on Macintosh workstations, that prevents users from copying the file.

See also Attributes.

Country object

A container object that represents a country where your network resides and organizes other Directory objects.

See also Object; Container object.

Create right

A file system right that grants the right to create new files or subdirectories, or to salvage a file after it has been deleted. Also, an object right that grants the right to create a new object in the Directory tree.

See also Rights.

Custom Device Module

(CDM) The driver component in NPA™ used to drive specific storage devices attached to the host adapter.

Cylinder

Used to organize data on a hard disk. Some hard disks have multiple platters with a read/write head for each side of each platter. The group of read/write heads moves in unison to the same track number on each platter side.This group of tracks is called a cylinder and is numbered according to the tracks it references. For example, the 0 cylinder is a group of 0 tracks.

Daemon

A process running in the background that can spawn (initialize) other processes with little or no user input. Daemons provide services such as printing, remote printing, and server advertising. Some daemon processes, such as the NetWare daemon, perform administrative functions and access the host file system.

Data fork

The part of the Macintosh file containing the series of bytes which makes up the actual file data. For example, the data fork of an application file contains the instructions that make the application work.

See also Macintosh files.

Data migration

The transfer of inactive data from a NetWare volume to an optical disc storage device. Data migration lets you move data to an optical disc storage device, called a *jukebox*; while NetWare still sees the data as residing on the volume. This frees valuable hard disk space for often-used files while still allowing access to infrequently used files.

See also Attributes; High Capacity Storage System.

Data protection

A means of ensuring that data on the network is safe. NetWare protects data primarily by maintaining duplicate file directories and by redirecting data from bad blocks to reliable blocks on the NetWare server's hard disk.

See also Directory Entry Table; Disk duplexing; Disk mirroring; System Fault Tolerance.

Data set

A set of data that can be manipulated by SBACKUP. Data sets can contain different items depending on which Target Service Agent they are related to.

See also Restore.

Default drive

The drive a workstation is currently using. The drive prompt, such as A:> or F:>, identifies the drive.

Default server

The server you attach to when you load the NetWare Requester™. The default server is the preferred server specified in your NET.CFG file.

Delete Inhibit (Di) attribute

A file system attribute that prevents any user from erasing the directory or file.

See also Attributes.

Delete right

An object right that grants the right to delete an object from the Directory tree.

See also Rights.

Delimiter

A symbol or character that signals the beginning or end of a command or of a parameter within a command.

Delimiters used in NetWare include the space, the comma (,), the period (.), the slash (/), the backslash (\), the hyphen (–), and the colon (:).

Destination server

The NetWare 4 server to which you migrate data files, bindery files, and other information from a previous NetWare version or another network operating system when upgrading to NetWare 4.

See also Source server.

DET

(Directory Entry Table) A table that contains basic information about files, directories, directory trustees, or other entities on the volume.

See also Directory entry.

Device driver

A program (usually an NLM) that controls devices (such as a printer, network board, diskette drive, hard disk, or monitor) attached to the computer. Device drivers expand an operating system's ability to work with peripherals because they control the software routines that make peripherals work.

See also LAN driver.

Device numbering

A method of identifying a device, such as a hard disk, so that the device can work on the network. Devices are identified by three sets of numbers:

- Physical address
- Device code
- Logical device number

Device sharing

The shared use of centrally located devices (such as printers, modems, and disk storage space) by users or software programs.

By attaching a device to a network that several workstations are logged in to, you can use resources more efficiently.

Directory

Directory: A common name for the NetWare Directory database, which organizes NDS objects in a hierarchical tree structure called the *Directory tree.* (*See also* NetWare Directory Services.)

directory: A component in the NetWare file system, used to contain files and subdirectories. (*See also* File system.)

Directory and file rights

Rights that control what a trustee can do with a directory or file.

See also Rights.

Directory caching

A method of decreasing the time used to determine a file's location on a disk.

See also Cache memory.

Directory database

A common name for the NetWare Directory database.

See also NetWare Directory database.

Directory entry

Basic information for NetWare server directories and files, such as

- File or directory name
- Owner
- Date and time of the last update (for files)
- Location of the first block of data on the network hard disk

Directory entries are located in a directory table on a network hard disk and contain information about all files on the volume. The server uses directory entries to track file location, changes made to the file, and other related file properties.

See also Directory Entry Table.

Directory Entry Table

(DET) A table that contains basic information about files, directories, directory trustees, or other entities on the volume.

See also Directory entry.

Directory hashing

A process that facilitates access to a file in a large volume by calculating the file's address both in cache memory and on the hard disk. When a workstation wants to read a file from the NetWare server, the server performs a hash algorithm that predicts an address on a hash table.

See also Cache memory.

Directory Map object

A leaf object that refers to a directory on a volume.

See also Object.

Directory path

The full specification that includes the server name, volume name, and name of each directory leading to the file system directory you need to access.

See also Drive mapping; File system.

Directory partition

A common name used for NetWare Directory partitions.

See also NetWare Directory partition.

Directory replica

A common name used for NetWare Directory replicas.

See also NetWare Directory replica.

Directory rights

Rights that control what a trustee can do with a directory.

See also Rights.

Directory services

A database built into NetWare 4 that maintains information about every resource on the network.

See also NetWare Directory Services.

Directory structure

A hierarchical structure that represents how Directory partitions are related to each other in the Directory database. (*See also* Directory tree.) Also, formerly used to describe the filing system of volumes, directories, and files that the NetWare server uses to organize data on its hard disks. (*See also* File system.)

Directory tree

A hierarchical structure of objects in the Directory database. The Directory tree includes container objects that are used to organize the network. The structure of the Directory tree can be based on a logical or physical organization of objects.

See also Browsing; Object; NetWare Directory partition; NetWare Directory replica.

Disk

A magnetically encoded storage medium in the form of a plate (also called a *platter*). The following types of disks are used with personal computers:

- **Hard disk**. Uses a metallic base and is usually installed within a computer or disk subsystem. (In some cases, they are removable.)
- **Floppy disk** (Also called a *diskette*). Uses a mylar base and is removable.
- **CD-ROM** (Compact Disc Read Only Memory). A small plastic optical disc that cannot be written to or erased.
- **Optical disc.** Either erasable and writable, or WORM (Write Once, Read Many).

See also Data protection; Disk partition; Hard disk.

Disk controller

An adapter, board, or chip set on the system board. This device controls how data is written to and retrieved from the disk drive. The disk controller sends signals to the disk drive's logic board to regulate the movement of the head as it reads data from or writes data to the disk.

See also Host Bus Adapter.

Disk driver

An NLM that forms the interface between the NetWare operating system and the hard disks. The disk driver talks to an adapter that is connected by an internal cable to the disk drives. Depending on the type of disk controller or adapter, one or more disk drives can be connected. Drivers can be loaded into the operating system during installation or at the command line.

Disk duplexing

A means of duplicating data to provide data protection. Disk duplexing consists of copying data onto two hard disks, each on a separate disk channel. Disk duplexing allows the same data to be written to all disks simultaneously. This protects data against the failure of a hard disk, or of the hard disk channel between the disk and the NetWare server. (The hard disk channel includes the disk controller and interface cable.)

Disk duplexing also allows *split seeks*. This sends read requests to whichever disk can respond first. Multiple read requests are also split between the duplexed disks for simultaneous processing.

See also Device numbering; Disk mirroring.

Disk format

The way in which a hard disk is prepared or structured so that it can receive data from the computer's operating system. Disk formatting is a function of the operating system.

Disk mirroring

The duplication of data from the NetWare partition on one hard disk to the NetWare partition on another hard disk. When you mirror disks, two or more hard disks on the *same channel* are paired. Blocks of data written to the original (primary) disk are also written to the duplicate (secondary) disk. The disks operate in tandem, constantly storing and updating the same files. Should one of the disks fail, the other disk can continue to operate without data loss or interruption. A problem in the channel would cause a failure in both disks.

See also Device numbering; Disk duplexing.

Disk partition

A logical unit that NetWare server hard disks can be divided into. In NetWare 4, a NetWare partition is created on each hard disk. Volumes are created from the pool of NetWare partitions.

See also Data protection; Hot Fix; Volume.

Disk subsystem

An external unit that attaches to the NetWare server and contains hard disks, a tape drive, optical drives, or any combination of these. The disk subsystem gives the server more storage capacity.

Distance vector

An algorithm that disseminates routing information to routers on a network. A router using the distance vector algorithm maintains only enough information to know how to reach the next router destination (hop) on the network. Distance vector routers periodically forward this information to each other, even if the information has not changed since the last update. Such broadcasts create unnecessary traffic on the network and consume router CPU time.

Distribution List object

A leaf object that represents a list of mail recipients. A member of a Distribution List can be a user, another Distribution List object, a Group object, or an Organizational Role object.

Domains

Memory segments in NetWare that allow you to separate NLM files from the operating system. NetWare 4.1 has two domains: OS domain and OS_PROTECTED domain. By loading an NLM in the OS_PROTECTED domain, you prevent a misbehaved NLM from writing to memory it should not have access to. This keeps a misbehaved NLM from bringing down the server.

See also Memory protection; Paging.

Don't Compress (Dc) attribute

A file system attribute that prevents files from being compressed.

See also Attributes.

Don't Migrate (Dm) attribute

A file system attribute that prevents files from being migrated to a secondary storage device (such as a tape drive or optical disc).

See also Attributes.

Don't Suballocate (Ds) attribute

A file system attribute that prevents an individual file from being suballocated, even if suballocation is enabled for the system. Use for files that are often enlarged or appended, such as certain database files.

See also Attributes.

DOS client

A workstation that boots with DOS and gains access to the network through one of the following:

- The NetWare DOS Requester™ software (for NetWare 4)
- A NetWare shell (for NetWare 2 and NetWare 3)

With DOS client software, users can perform networking tasks. These tasks include mapping drives, capturing printer ports, and sending messages. Using the NetWare Requester in NetWare 4, users can change contexts.

See also Client.

DOS device

A storage unit compatible with the DOS disk format—usually a hard disk or tape backup unit.

DOS Requester, NetWare

The DOS client software portion of NetWare 4.

See also NetWare DOS Requester.

DOS setup routine

The routine that sets up the system configuration of your DOS client or NetWare server. The setup routine records the system's built-in features (add-on boards, hard drives, disk drives, ports, math coprocessor) and available system memory. It also lets you set date and time, password, and keyboard speed.

DOS version

The version number and name of the DOS you are using (such as Novell DOS™ or MS-DOS). You must create a DOS directory for each workstation type or DOS version you use and load the DOS files into it.

See also File system; Login scripts.

Drive

Physical drive. A storage device, such as a disk drive or tape drive, that data is written to and read from. A drive that is physically contained in or attached to a workstation is called a *local drive*.

Logical drive. An identification for a specific directory located on a disk drive. For example, a network drive points to a directory on the network, rather than to a local disk.

Drive mapping

A pointer to a location in the file system, represented as a letter assigned to a directory path on a volume. To locate a file, you follow a *path* that includes the volume, directory, and any subdirectories leading to the file.

Search drive mappings enable the system to locate a program even if it is not located in the directory you are working in.

See also Search drive.

Driver

Software that forms the interface between the NetWare operating system and devices such as hard disks or network boards.

See also Device driver; LAN driver.

Dual Processing

A NetWare® SFT III™ configuration that assigns parts of the operating system to separate processors. Because SFT III is split into two engines (the IOEngine and the MSEngine), each engine can be run on a separate CPU, creating a dual processing system.

Duplexing

A means of duplicating data to provide data protection. Disk duplexing consists of copying data onto two hard disks, each on a separate disk channel.

See also Disk duplexing.

Dynamic configuration

A means of allowing the NetWare server to allocate resources according to need and availability. When the server boots, all free memory is assigned to file caching. As demand increases for other resources (directory cache buffers, for example), the number of available file cache buffers decreases.

Dynamic memory

The most common form of memory, used for RAM. Dynamic memory requires a continual rewriting of all stored information to preserve data. If dynamic memory is too slow for a computer's microprocessor, overall performance suffers while the CPU waits for requested information to arrive from memory.

A continuous electrical current is necessary to maintain dynamic memory. All data is lost from dynamic memory when the power is turned off.

Effective rights

The rights that an object can actually exercise to see or modify a particular directory, file, or object. An object's effective rights to a directory, file, or object are calculated by NetWare each time that object attempts an action.

Effective rights to a file or directory are determined by

- An object's trustee assignments to the directory, file, or other object
- Inherited rights from an object's trustee assignments to parent directories
- Trustee assignments of Group objects that a User object belongs to
- Trustee assignments of objects listed in a User object's "Security Equal To" list

See also Inherited Rights Filter; Security; Trustee.

EGP

(Exterior Gateway Protocol) A protocol that exchanges network reachability information between autonomous systems. EGP is part of the TCP/IP protocol suite.

See also Autonomous system.

Elevator seeking

Organizes the way data is read from hard disk storage devices. A shared network disk drive and its related channel can quickly become clogged with disk I/O requests. Elevator seeking logically organizes disk operations as they arrive at the server for processing. Elevator seeking improves disk channel performance by significantly reducing disk head thrashing (rapid back-and-forth movements) and by minimizing head seek times.

Embedded SCSI

A hard disk that has a SCSI and a hard disk controller built into the hard disk unit.

See also SCSI.

Erase right

A file system right that grants the right to delete directories, subdirectories, or files.

See also Rights.

Ethernet configuration

The setup that allows communication using an Ethernet environment. In an Ethernet environment, stations communicate with each other by sending data in frames along an Ethernet cabling system. Different Ethernet standards use different frame formats. NetWare 4 uses the IEEE 802.2 standard by default.

See also Link Support Layer; Multiple Link Interface Driver; Open Data-Link Interface; Packet.

Execute Only (X) attribute

A file system attribute that prevents a file from being copied.

See also Attributes.

Extended AppleTalk network

An AppleTalk network that is capable of supporting the AppleTalk Phase 2 extensions, such as zone lists and network ranges.

See also AppleTalk Phase 2.

Expansion bus

A common pathway between hardware devices. A bus connects the CPU to its main memory and the memory that resides on the control units of the peripheral devices. A bus allows you to transfer data from the system board to plug-in peripherals. An expansion bus allows you to expand your PC system using adapter cards by providing a data path from the adapter to the CPU.

See also Local bus.

External Entity object

A leaf object that represents a non-native NDS object that has been imported into NDS or registered in NDS.

NetWare Message Handling Service™ (NetWare MHS™) uses External Entity objects to represent users from non-NDS environments and provides an integrated address book for sending mail.

See also NetWare MHS Services.

Failure handling

The SFT III process that prevents system downtime caused by a single hardware failure in a NetWare server. With SFT III, if one NetWare server fails, its mirrored partner automatically takes over. This failure-handling process is instantaneous, automatic, and transparent to NetWare workstations.

Fake root

A subdirectory that functions as a root directory, allowing you to assign users rights at the subdirectory level. Some applications cannot be run from subdirectories; they read files from and write files to the root directory. However, for security, you should not assign users rights at the root or volume directory level. NetWare allows you to map a drive to a fake root (a directory where rights can be assigned to users).

See also Security.

FAT

(File Allocation Table) An index table that points to the disk areas where a file is located.

See also File Allocation Table.

Fault tolerance

A means of protecting data by reconciling bad disk blocks or by providing data duplication. (*See also* System Fault Tolerance.)

Also, distributing the Directory database among several servers to provide continued authentication and access to object information should a server go down. (*See also* NetWare Directory replica.)

File Allocation Table

(FAT) An index table that points to the disk areas where a file is located. Because one file may be in any number of blocks spread over the disk, the FAT links the file together.

See also Turbo FAT index.

File caching

The use of NetWare server RAM to improve file access time.

See also Cache memory.

File compression

A means of allowing more data to be stored on server hard disks by compressing (packing) files that are not being used. File compression is managed internally by the NetWare operating system. Users can flag their files or directories so they are compressed after being used, or flag them so they are never compressed.

After compression is enabled, files flagged Immediate Compress (Ic) are compressed immediately; other files are automatically compressed when they have not been accessed for a specific amount of time. Files are decompressed when a user accesses them again.

If a disk error or a power failure occurs during compression, the original file is retained.

See also Attributes.

File handle

A number used to refer to or identify a file.

File indexing

The method of indexing FAT entries for faster access to large files.

See also File Allocation Table.

File locking

The means of ensuring that a file is updated correctly before another user, application, or process can access the file.

See also Record locking; Semaphore.

File rights

Rights that control what a trustee can do with a file.

See also Rights.

File Scan right

A file system right that grants the right to see the directory and file with the DIR or NDIR directory command.

See also Rights.

File server

A name used in previous NetWare versions for the computer that runs NetWare operating system software; now referred to as the *NetWare server.*

See also NetWare server.

File sharing

A feature of networking that allows more than one user to access the same file at the same time.

See also Security.

File system

(Formerly *directory structure*) The system the NetWare server uses to organize data on its hard disks. Each file is given a filename and stored at a specific location in a hierarchical filing system so that files can be located quickly.

The NetWare server is divided into one or more volumes. Volumes are divided into directories containing files or subdirectories. A directory can contain any number of files and subdirectories.

See also Directory; Directory Entry Table; Drive mapping; Security.

File Transfer Protocol

(FTP) A set of control procedures used to prevent errors in information transmitted between network stations. FTP is part of the TCP/IP protocol suite. The data is sent from one station to another in packets. Each packet includes a discrete number that is derived from the data that makes up the packet, according to a mathematical algorithm. The algorithm is applied to each data packet a second time when it arrives on the receiving end. If the number on the receiving end does not match the number included in the packet, the receiving station sends a signal to the transmitting station requesting that the packet be resent.

Filename extension

The extension used after the period in filenames. Under the FAT system used by DOS and OS/2, filename extensions can be up to three characters in length. Under the High Performance File System (HPFS) used by OS/2, filename extensions are not restricted to three characters.

Foreign E-mail address

Specifies a mailbox that resides in a foreign E-mail system. An object can have only one foreign E-mail address.

Foreign E-mail alias

Specifies an object's aliases as known in a foreign messaging system. A foreign E-mail alias is the return address value used when the NetWare MHS user sends E-mail to an X.400 user. An object can have several foreign E-mail aliases, one for each type of foreign E-mail system.

Form

In a NetWare printer command, the name and size of the paper used for a print job.

See also Printer form.

Frame

A packet data format for a given media. Some media, such as Ethernet 802.2, Ethernet 802.3, Ethernet II, Ethernet SNAP, token ring, or token ring SNAP, support multiple packet formats (frames). For NetWare 4, the default Ethernet frame type is 802.2.

See also Ethernet configuration.

FTP

(File Transfer Protocol) A set of control procedures used to prevent errors in information transmitted between workstations. FTP is part of the TCP/IP protocol suite.

See also File Transfer Protocol.

Gateway

A link between two networks. A gateway allows communication between dissimilar protocols (such as NetWare and non-NetWare networks) using industry standard protocols such as TCP/IP, X.25, or SNA.

Group object

A leaf object that represents a grouping of several User objects; used to provide collective, rather than individual, network administration. You can create Group objects based on who uses the same applications, printers, or print queues; who performs similar tasks; or who has similar needs for information. You can also use Group objects to simplify trustee assignments.

See also Effective rights; Object.

Handle

A pointer used by a computer to identify a resource or feature. For example, a directory handle identifies a volume and a directory, such as SYS:PUBLIC. Other types of handles used to access NetWare 4 include file handles, video handles, request handles, device handles, and volume handles.

Handshaking

The initial exchange between two data communication systems prior to and during data transmission to ensure proper data transmission. A handshake method (such as XON/XOFF) is part of the complete transmission protocol. A serial (asynchronous) transmission protocol might include the handshake method (XON/XOFF), baud rate, parity setting, number of data bits, and number of stop bits.

See also Serial communication.

Hard disk

A high-capacity magnetic storage device that allows a user to write and read data. Hard disks can be network or local workstation disks.

See also Data protection; Disk driver; Host Bus Adapter; Partition (disk).

Hashing

A process that facilitates access to a file in a large volume by calculating the file's address both in cache memory and on the hard disk.

See also Cache memory; Directory hashing.

HBA

(Host Bus Adapter) A board that acts as an interface between the host microprocessor and the disk controller.

HCSS

(High Capacity Storage System) A system that increases the data storage capacity of a NetWare server by integrating an optical disc library, or jukebox, into the NetWare file system.

See also Data migration.

Hexadecimal

A base-16 numeric notation system used to specify addresses in computer memory. In hexadecimal notation, the decimal numbers 0 through 15 are represented by the decimal digits 0 through 9 and the alphabetic characters A through F (A = decimal 10, B = decimal 11, and so on).

Hidden (H) attribute

A DOS and OS/2 attribute that hides a directory or file from the DOS or OS/2 DIR command and prevents the directory or file from being deleted or copied.

See also Attributes.

High Capacity Storage System

(HCSS) A system that increases the data storage capacity of a NetWare server by integrating an optical disc library, or jukebox, into the NetWare file system.

See also Data migration.

Home directory

A private network directory that the network supervisor can create for a user. Each user's login script should contain a drive mapping to the user's home directory.

Hop count

The number of cable segments a message packet passes through on the way to its destination on a network or internetwork. The destination network can be no more than 16 hops from the source.

See also Partition (disk).

Host

A NetWare server from which you run SBACKUP. A storage device and a storage device controller are attached to it.

See also Backup; Target.

Host Adapter Module

(HAM) The driver component used to drive specific host adapter hardware in NPA.

Host Bus Adapter

(HBA) A SCSI adapter board or disk controller that adds a bus through which peripheral devices (such as hard disks, tape drives, and CD-ROM drives) are connected to the computer. These devices typically have embedded controllers.

See also Hard disk; SCSI bus.

Hot Fix

A method NetWare uses to ensure that data is stored safely. Data blocks are redirected from faulty blocks on the server's disk to a small portion of disk space set aside as the *Hot Fix*™ *Redirection Area*. Once the operating system records the address of the defective block in a section of the Hot Fix area reserved for that purpose, the server does not attempt to store data in defective blocks. By default, 2% of a disk partition's space is set aside as the Hot Fix Redirection Area.

Hot Fix is always active unless the disk fails and is inoperative or the redirection area is full. Hot Fix, together with read-after-write verification, enables a hard disk to maintain data integrity.

See also Data protection.

Hub

A device that modifies transmission signals, allowing the network to be lengthened or expanded with additional workstations.

Two kinds of hubs exist:

- **Active hubs.** Used to amplify transmission signals in network topologies. Use active hubs to add workstations to a network or to extend the cable distance between stations and the server.
- **Passive hubs.** Used in certain network topologies to split a transmission signal, allowing additional workstations to be added. Since a passive hub cannot amplify the signal, it must be cabled directly to a station or to an active hub.

ICMP

(Internet Control Message Protocol) A protocol in the TCP/IP suite that sends packets containing information about network failures, such as inoperative nodes and gateways or congestion at a gateway.

See also Internet Control Message Protocol.

IDE

A standard interface for a hard disk.

See also Integrated Drive Electronics.

Identifier variables

Variables used in login scripts that allow you to enter a variable (such as LOGIN_NAME) in a login script command, rather than specific information (such as RICHARD).

See also Login scripts.

Immediate Compress (Ic) attritube

A file system attribute that causes files to be compressed as soon as the operating system can do so. The operating system does not wait for a specific event (such as a time delay) before compressing the file.

See also Attributes.

Indexed (I) attribute

A status flag set automatically when a file exceeds a set size, indicating that the file is indexed for fast access.

See also Attributes.

Inherited Rights Filter

(IRF) A list of rights created for every file, directory, and object. The IRF controls the rights that a trustee can inherit from parent directories and container objects.

By default, the IRF allows every right to be inherited from the parent directory or container object. The IRF cannot grant rights; it can only allow or revoke rights. To allow a right, the right must exist in the parent directory or container object and the IRF. To revoke a right, the right must exist in the parent directory or container object and then be removed from the IRF.

The IRF is ignored whenever a trustee has an explicit trustee assignment to that file, directory, or object.

See also Effective rights; Security.

Integrated Drive Electronics

(IDE) A hard disk drive standard interface. The IDE integrates controller electronics onto the drive. The controller connects to a paddleboard that may be external to, or on, the system board. The paddleboard then interfaces with the bus to the CPU. You can identify an IDE bus by its 40-pin connector, as opposed to a SCSI bus, which has a 50-pin connector.

See also Hard disk.

Internal network number

A logical network number that identifies an individual NetWare 3 or NetWare 4 server. On IPX networks, the internal network number must also be different from the IPX external network number.

See also IPX external network number; IPX internal network number.

International use of NetWare 4.1

The adaptation of NetWare 4 for use with multiple languages. The NetWare 4 operating system, NLM programs, and utilities use English as the default language but can be set to several other languages.

Internet Protocol

(IP) The network layer protocol of the TCP/IP suite of protocols. IP enables dissimilar nodes in a heterogenous environment to communicate with one another.

See also TCP/IP.

Internetwork

Two or more networks connected by a router. Each network has a unique network number. Users on an internetwork can use the resources (files, printers, hard disks) of all connected networks, provided the users have security clearance.

See also IPX external network number; Router.

Interoperability

The ability to use products from different vendors within the same system. Communication protocols, such as IP or AFP, can be used in ODI to process information from the network without the user having to know each protocol's required method of packet transmission.

Interoperability also means that an application running on different platforms (such as Macintosh or UNIX) can share files.

See also Open Data-Link Interface.

Interrupt mode

A printer configuration option through which the data port sends a signal, or interrupt, to the port driver (NPRINTER) when it is ready to accept another character. The interrupt instructs the CPU to suspend its other processing activities to service the needs of the port in question.

See also Polled mode.

IP address

Identifies the network to which the host server is attached. The address is normally in four segments, each separated by a period (for example, 87.34.53.12). Individual numbers must be between 0 and 255.

IPX

(Internetwork Packet Exchange) A Novell communication protocol that sends data packets to requested destinations (such as workstations or servers). IPX allows packet addressing within a single network or in an internetwork environment (that is, two or more networks connected by a router, where each network has a unique IPX external network number). Through IPX, incoming data packets are directed to the proper area within the operating system of the workstation or NetWare server.

See also Communication protocols; Internetwork; IPXODI; LAN driver; NETX; Open Data-Link Interface.

IPX external network number

A network number that uniquely identifies a network cable segment. An IPX external network number is a hexadecimal number, one to eight digits (1 to FFFFFFFE). The number is arbitrary and is assigned when the IPX protocol is bound to a network board in the server. You can bind IPX with multiple frame types to the same network board.

The terms *network number* and *network address* are sometimes used to refer to the IPX external network number.

See also IPX internal network number; IPX internetwork address; Network numbering.

IPX internal network number

A logical network number that identifies an individual NetWare server. The IPX internal network number is a hexadecimal number, one to eight digits (1 to FFFFFFFE), that is assigned to the server during installation. Each server on a network must have a unique IPX internal network number. The IPX internal network number of any node must also be different from any IPX external network number on the internetwork.

See also AUTOEXEC.NCF; IPX external network number.

IPX internetwork address

A 12-byte number (represented by 24 hexadecimal characters) divided into three parts. The first part is the 4-byte (8-character) IPX external network number. The second part is the 6-byte (12-character) node number. The third part is the 2-byte (4-character) socket number.

See also IPX external network number; Node number; Socket.

IPXODI

(Internetwork Packet Exchange Open Data-Link Interface) A module that takes the workstation requests the NetWare DOS Requester has determined are for the network, packages them with transmission information (such as their destination), and transfers them to the LSL. IPXODI requires each packet to have an initialized header. The header specifies information that targets network delivery, announcing where the packet came from, where it is going, and what happens after delivery.

See also IPX; Link Support Layer; NetWare DOS Requester; SPX.

Jukebox

A high-capacity storage device, sometimes called an *optical disc library,* that uses an autochanger mechanism to mount and dismount optical discs as they are needed. A jukebox typically contains one to four optical disc drives. A picker rotates, flips, and transports disks to and from the storage slots, drives, and mail slot. The mail slot is the location in the jukebox used to insert and remove the optical disc cartridge.

See also High Capacity Storage System.

Jumper block

A group of jumper pins that can be connected (*jumpered*) or left unconnected to make hardware configuration settings on a circuit board.

LAN

(LAN) A network located within a small area or common environment, such as in a building or a building complex.

See also Wide area network.

LAN driver

An NLM that interfaces with a network board. A LAN driver serves as a link between a station's operating system and the physical network components.

See also Link Support Layer; Multiple Link Interface Driver; NetWare Loadable Module; Open Data-Link Interface.

Large Internet Packet

(LIP) Functionality that allows the maximum size of internetwork packets to be increased. (Formerly, the maximum size was 576 bytes.) In NetWare 4, LIP allows the workstation to determine the packet size based on the maximum size supported by the router. LIP functionality is implemented for DOS clients through the station's NET.CFG file.

Leaf objects

Objects that do not contain any other objects and are located at the end of a branch in the Directory tree.

See also Object.

Link state

A routing algorithm that builds and maintains a logical map of the entire network. A link state router accomplishes this by sending a packet containing information about all its links—connections to networks and other routers—to all other link state routers on the network. This process is known as *flooding*. Each router uses this information to build the network map.

When each link state router has the same view (map) of the network, the network is said to have *converged*. Link state routers multicast their link information only when a change occurs in a route or service.

See also NetWare Link Services Protocol.

Link Support Layer

(LSL) An implementation of the ODI specification that serves as an intermediary between the NetWare server's LAN drivers and the communication protocols, such as IPX, AFP, or TCP/IP. The LSL allows one network board to service several communications protocol stacks. It also allows several network boards to service the same protocol stack.

See also Open Data-Link Interface.

LIP

(Large Internet Packet) A functionality that allows the maximum size of internetwork packets to be increased.

See also Large Internet Packet.

Loadable module

A program you can load and unload from a server or a workstation while the attendant operating system is running. Two common types are NLM programs and Virtual Loadable Module™ (VLM) programs.

See NetWare Loadable Module; Virtual Loadable Module.

Loading and unloading

The process of linking and unlinking NLM programs to the NetWare operating system. NLM programs can be loaded and unloaded while NetWare is running.

See also NetWare Loadable Module.

Local area network

(LAN) A network located within a small area or common environment, such as in a building or a building complex.

See also Wide area network.

Local bus

A 32-bit internal path that connects the CPU directly to memory, video, and disk controllers. The local bus allows data transfer to and from the CPU to memory and peripherals at CPU speeds (such as 33,66, and 90 MHz).

Local bus architecture is utilized by Peripheral Component Interconnect (PCI) and Video Electronics Standards Association (VESA) buses.

See also Expansion bus.

Local drive

A common name for a *physical drive* attached to a workstation.

See Drive.

Logical memory

Memory that may not have contiguous addresses, but which appears contiguous to NetWare 4 processes.

See also Paging.

Login

The procedure that provides access to the network by using the LOGIN command. When a user initiates a login request, the operating system looks for security rights; the user is then asked for a password. All security information is placed in the NetWare server's connection list and the user is said to be logged in. At this point, LOGIN executes one or more login scripts (which initialize environment variables, map network drives, and so on).

See also LOGIN directory; Login restrictions; Login scripts; logout.

LOGIN directory

The SYS:LOGIN directory, created during network installation; contains the LOGIN and NLIST utilities. Users can use these utilities to log in and view a list of available NetWare servers. For NetWare users running OS/2, the corresponding LOGIN directory is SYS:LOGIN/OS2.

See also File system; MAIL directory; PUBLIC directory; SYSTEM directory.

Login restrictions

Limitations on user accounts that control access to the network, such as

- Requiring a password
- Setting account limits
- Limiting disk space
- Specifying the number of connections
- Setting time restrictions

When a user violates login restrictions, NetWare disables the account and no one can log in using that username. This prevents unauthorized users from logging in.

Login scripts

Files containing commands that set up users' workstation environments whenever they log in. Login scripts are similar to batch files and are executed by the LOGIN utility.

You can use login scripts to

- Map drives and search drives to directories
- Display messages
- Set environment variables
- Execute programs or menus

See also Drive mapping.

Logout

A procedure that breaks the network connection and deletes drives mapped to the network. If you log out without specifying a NetWare server name in the LOGOUT command, the station connections and drives mapped to all servers are terminated. To log out from one server and remain attached to the other servers, specify the server name in the LOGOUT command.

Long machine type

A six-letter name representing a DOS workstation brand. Use the long machine type in container login scripts (using the MACHINE identifier variable) to automatically map a drive to the correct version of DOS assigned to the station.

See also DOS version; Login scripts; Short machine type.

LPT ports

The parallel printer ports of a personal computer.

See also Parallel port.

LSL

(Link Support Layer) An intermediary between the NetWare server's LAN drivers and communication protocols, such as IPX, AFP, or TCP/IP.

See also Link Support Layer.

Macintosh client

A Macintosh computer that attaches to the network. The Macintosh client can store data on, and retrieve data from, a NetWare server running NetWare for Macintosh name space support modules. The Macintosh client can also run executable Macintosh network files, share files with other clients (DOS, MS Windows, OS/2, and UNIX), and monitor queues.

See also Client.

Macintosh files

Files used on Macintosh computers. A Macintosh file contains two parts: the data fork and the resource fork.

- **Data fork.** Contains information (data) specified by the user.
- **Resource fork.** Contains file resources, including Macintosh-specific information such as the windows and icons used with the file.

To store Macintosh files on a NetWare server, the Macintosh name space module must be linked with the NetWare operating system container.

See also Name space support.

Mailbox ID

A unique name that specifies the directory in which all of the object's inbound mail is placed.

Mailbox location

The name of the messaging server where an object's mailbox resides.

MAIL directory

The SYS:MAIL directory, created during network installation and used by mail programs that are compatible with NetWare.

See also File system; LOGIN directory; PUBLIC directory; SYSTEM directory.

Major resource

A category of data defined by the Target Service Agent and recognized by SBACKUP. A major resource contains data that can be backed up as a group, such as the data on a server or volume.

See also Backup; Minor resource; Target Service Agent; Transaction Tracking System.

Map

For DOS and OS/2 clients, to assign a drive letter to a directory path on a volume. For example, if you map drive F: to the directory SYS:ACCTS/RECEIVE, you access that directory every time you change to drive F:.

See also Drive mapping.

Master replica

The Directory replica used to create a new Directory partition in the Directory database or to read and update Directory information. Although many Directory replicas can exist in the Directory, only one can be the master replica. The master replica is always considered to be the most accurate Directory replica.

See also NetWare Directory replica.

Memory

The internal dynamic storage of a computer that can be addressed by the computer's operating system; referred to frequently as RAM (random access memory). Memory accepts and holds binary data. To be effective, a computer must store the data that is operated on, as well as the program that directs the operations to be performed. Memory stores information and rapidly accesses any part of the information upon request.

Memory allocation

The process of reserving specific memory locations in RAM for processes, instructions, and data. When a computer system is installed, the installer may allocate memory for items such as disk caches, RAM disks, extended memory, and expanded memory. Operating systems and application programs allocate memory to meet their requirements, but they can use only that memory actually available to them.

Memory board

An add-on board that increases the amount of RAM within a personal computer.

See also Memory; RAM.

Memory protection

The structuring of memory resources in NetWare 4 that guards NetWare server memory from corruption by NLM programs. In NetWare 4, memory can be divided up into two domains: OS domain and OS_PROTECTED domain.

By default, the operating system and NLM programs run in the OS domain. You can load the DOMAIN.NLM to create the OS_PROTECTED domain. Some NLM programs can be loaded in the OS_PROTECTED domain. NLM programs running in the OS_Protected domain have limited entry points to the OS domain, thus preventing unauthorized memory access.

The OS_PROTECTED domain allows you to test your NLM programs without risking server memory corruption. Once you load an NLM in the OS_PROTECTED domain and find it safe, you can load it into the OS domain, where it can run most efficiently.

See also Domains; Paging.

Message packet

A unit of information used in network communication.

See also Packet.

Message Routing Group object

A leaf object that represents a group of messaging servers that can send and receive messages among themselves.

See also Messaging Server object.

Message system

A communications protocol that runs on top of IPX. It provides an engine that allows a node on the network to send messages to other nodes. A set of APIs (application program interfaces) gives programs access to the message system.

See also IPX.

Messaging Server object

A leaf object that represents a messaging server which resides on a NetWare server. A Messaging Server object is automatically created in the Directory tree when you install NetWare MHS on a NetWare server. The Messaging Server object is automatically placed in the same context as the NetWare Server object.

A messaging server can service an unlimited number of mailboxes. The amount of disk space available for mailboxes is the only limiting factor.

Migrated (M) attribute

A status flag, set automatically, that indicates a file has been migrated.

See also Attributes; Data migration.

Migration (operating system)

The conversion of servers from NetWare 2 or NetWare 3, or from another network operating system, to NetWare 4. (Do not confuse migration from one version of NetWare to another with data migration, which refers to moving files to near-line or offline storage devices.

See also Upgrade.

Migration (protocol)

The conversion of a server, router, or network from IPX to NetWare Link Services Protocol (NLSP), or from TCP/IP to Open Shortest Path First (OSPF) protocol.

See also IPX; NetWare Link Services Protocol.

Minor resource

A category of data defined by the Target Service Agent and recognized by SBACKUP. A minor resource might be located in the directory structure below the selected major resource (for example, directories, subdirectories, or files).

See also Backup; Major resource; Target Service Agent; Transaction Tracking System.

Mirroring

The duplication of data from the NetWare partition on one hard disk to the NetWare partition on another hard disk.

See also Disk mirroring.

MLID

(Multiple Link Interface Driver) A device driver written to the ODI specification that handles the sending and receiving of packets to and from a physical or logical LAN medium.

See also Multiple Link Interface Driver; Open Data-Link Interface.

Modify bit

A file attribute set by the operating system when a file is changed; indicates that data has been modified. The NetWare modify bit, called the Archive Needed attribute, appears as an "A" wherever file attributes are listed. When a backup is performed, SBACKUP checks to *see* whether modify bits are set and backs up only those files that have their modify bit set.

See also Backup.

Modify right

A directory or file right that grants the right to change the attributes or name of a directory or file.

See also Rights.

MS Windows client

A workstation that boots with DOS and gains access to the network through either

- The NetWare DOS Requester and its VLM programs (for NetWare 4)
- A NetWare shell (for NetWare versions earlier than NetWare 4)

The computer also runs MS Windows and, with the client software, can perform networking tasks in the MS Windows environment. These tasks include mapping drives, capturing printer ports, sending messages, and changing contexts.

See also Client.

Multiple-byte character

A single character made up of more than one byte. One byte allows 256 different characters. Since the number of ASCII characters equals 256, a computer can handle each ASCII character with one byte. Asian character sets, however, include more than 256 characters; in this case, a computer must use two bytes for each character.

Multiple Link Interface Driver

(MLID) A device driver written to the ODI specification that handles sending packets to and receiving packets from a physical or logical LAN medium.

See also Open Data-Link Interface.

Multiple name space support

Allows various workstations running different operating systems to create their own familiar naming conventions. With multiple name space support, each file stored on a given volume has a name that any workstation (running off of any operating system) can recognize. This name is stored in a file entry in the volume's DET.

Volumes that support multiple name spaces use one file and directory entry for each name space supported. The same applies to directory names. For example, a NetWare server configured to support DOS and Macintosh filenames would generate two 128-byte file entries for every file.

See also Name space support.

Multiserver network

A single network that has two or more NetWare servers. On a multiserver network, users can access files from any NetWare server they have access rights to. A multiserver network is not the same as an internetwork, where two or more networks are linked through a router.

See also Internetwork; Network numbering.

Name context

The position of an object in the Directory tree.

See also Context.

Name space support

NLM files that allow you to store non-DOS filenames on a NetWare 3 or 4 server. Files appear in native mode to users at different types of workstations. Name space NLM files have a .NAM extension. (For example, MAC.NAM and OS2.NAM.)

See also Multiple name space support; NetWare Loadable Module.

NCP

(NetWare Core Protocol) Procedures that a server's NetWare operating system follows to accept and respond to workstation requests.

See also NetWare Core Protocol.

NDS

(NetWare Directory Services) A relational database that is distributed across your entire network.

See also NetWare Directory Services.

NetBIOS

IBM's standard protocol for applications developed to run peer-to-peer communications on the IBM PC network and the token ring network. NetBIOS has become widely accepted as a standard for network interfacing.

The NetWare Client™ for DOS and MS Windows provides a NetBIOS driver that emulates the NetBIOS protocol. This emulator allows NetWare IPX to interface with the NetBIOS Interrupt 5Ch and an alternate interface, Interrupt 2Ah.

NET.CFG

A workstation boot file, similar to CONFIG.SYS in DOS, that contains configuration values that are read and interpreted when your workstation starts up. NET.CFG is created with an ASCII text editor and needs to be included on the workstation boot diskette with other boot files. NET.CFG replaces SHELL.CFG, used in earlier NetWare versions.

The configuration values in NET.CFG adjust the operating parameters of the NetWare DOS Requester, IPX, and other workstation software.

NetSync cluster

Includes one NetWare 4.1 server running NETSYNC4 and up to twelve NetWare 3.1*x* servers attached to it.

NetWare Client for DOS and MS Windows

Software that connects DOS and MS Windows workstations to NetWare networks and allows DOS and MS Windows users to share network resources.

NetWare Client for OS/2

Software that connects OS/2 workstations to NetWare networks and allows OS/2 users to share network resources.

NetWare Core Protocol

(NCP) Procedures that a server's NetWare operating system follows to accept and respond to workstation requests. NetWare Core Protocols exist for every service a station might request from a server. Common requests handled by NCP include creating or destroying a service connection, manipulating directories and files, opening semaphores, altering the Directory, and printing.

See also Communication protocols; IPX.

NetWare Directory database

The database (commonly referred to as *the Directory*) that organizes NetWare Directory Services objects in a hierarchical tree structure called the *Directory tree*.

See also NetWare Directory Services.

NetWare Directory partition

A logical division of the NetWare Directory database. A Directory partition forms a distinct unit of data in the Directory tree that you use to store and replicate Directory information. Each Directory partition consists of a container object, all objects contained in it, and data about those objects. Directory partitions do not include any information about the file system or the directories and files contained there.

See also NetWare Directory replica; NetWare Directory Services.

NetWare Directory replica

A copy of a NetWare Directory partition. For the Directory database to be distributed across a network, it must be stored on many servers. Rather than storing a copy of the whole Directory database on each server, Directory replicas of each Directory partition are stored on many servers throughout the network. You can create an unlimited number of Directory replicas for each Directory partition and store them on any server.

See also NetWare Directory partition.

NetWare Directory Services

(NDS) A relational database that is distributed across your entire network. NDS provides you with global access to all network resources to which you have been given rights, regardless of where they are physically located.

NDS treats all network resources as objects in a distributed database known as the *NetWare Directory database*, also referred to as the *Directory*.

NetWare Directory Services management request

A request that controls the physical distributeion of the Directory database. Through these requests, network adminstrators can create new Directory partitions and manage their Directory replicas.

See also NetWare Directory partition; NetWare Directory replica.

NetWare Directory Services request

A request made to the Directory database by users or network supervisors.

NetWare DOS Requester

A group of files that provide NetWare support for DOS and MS Windows client workstations. These files consist of a number of VLM™ programs and a single executable file (VLM.EXE) thatmanages operation of the .VLM files.

NetWare Link Services Protocol

(NLSP) A link state routing protocol designed by Novell for IPX internetworks. NLSP™ transfers routing information between routers and makes routing decisions based on that information. NLSP routers exchange link information such as network connectivity, path costs, IPX network numbers, media types, etc. By exchanging this information with its peer routers, each router builds and maintains a complete logical map of the network.

Unlike RIP and SAP, which periodically broadcast routing and service information, NLSP multicasts routing information only when a change occurs in a route or service on the network. To communicate with NetWare clients, NLSP routers use RIP.

See also Link State.

NetWare Loadable Module

(NLM) A program you can load and unload from server memory while the server is running. (Some NLM programs are loaded automatically because other NLM programs cannot run without them.) When loaded, an NLM program is dynamically linked to the operating system, and the NetWare server allocates a portion of memory to it. The NLM uses the memory to perform a task and then returns control of the memory to the operating system when the NLM is unloaded. When an NLM is unloaded, all allocated resources are returned to the operating system.

NetWare 4 has four types of NLM programs:

- Disk drivers (.DSK extension)
- LAN drivers (.LAN extension)
- Management utilities and server applications modules (.NLM extension)
- Name space support (.NAM extension)

Some NLM programs, such as utilities, can be loaded, used, and then unloaded. Other NLM programs, such as LAN driver and disk driver NLM programs, must be loaded every time the server is booted. NCF files (STARTUP.NCF and AUTOEXEC.NCF) allow you to store NLM commands that you want loaded every time the NetWare server is booted.

See also Loading and unloading; Memory protection.

NetWare managed node

A NetWare 4 server that has the NetWare Management Agent™ software enabled, making more information available to management console software than is available with IPX/SPX function calls.

NetWare MHS Services

Services that allow users to communicate electronically across a network. Using messaging services, users can perform tasks such as exchanging electronic mail, sharing calendars, and scheduling facilities. To provide messaging services, NetWare uses a messaging server, a Distribution List object, a Message Routing Group object, an External Entity object, and a Postmaster.

NetWare Name Service

(NNS) A naming service designed to provide more transparent access to resources in NetWare installations. NNS is no longer available or supported.

See also NetWare Directory Services.

NetWare Networked File System*

(NetWare NFS) Software that transparently integrates UNIX systems with NetWare 4 file systems and resources to give UNIX users access to the NetWare environment from their native operating system.

NetWare NFS*

Software that transparently integrates UNIX systems with NetWare 4 file systems.

See also NetWare Networked File System.

NetWare operating system

The network operating system developed by Novell, Inc. NetWare runs on the server and provides several functions to the network and the applications running on it, including

- File and record locking
- Security
- Print spooling
- Interprocess communications

The NetWare operating system also determines performance, multivendor support, and reliability of the network.

NetWare partition (disk)

A partition created on each network hard disk, from which NetWare volumes are created.

See also Disk partition.

NetWare protocols and transports

The components of NetWare software that allow client workstations to communicate and be understood on the network. A *protocol* manages data and a *transport* manages application messages. A protocol and transport can be provided by one piece of software or by many. In order for client workstations to communicate on the network, they must use a protocol identical to the one used on the network. However, workstations can be configured to use multiple protocols.

NetWare Runtime

A single-user version of the NetWare 4 operating system that provides NetWare services to clients of NLM programs.

NetWare server

A computer that runs NetWare operating system software. A NetWare server regulates communications among personal computers attached to it and to shared resources, such as printers. A NetWare 4.1 server must have at least one hard disk, either internal or external, and a recommended minimum 8 MB of RAM. The server must also contain at least one network board.

NetWare Server for OS/2

Device drivers that allow the NetWare 4 operating system to run as a nondedicated server on an OS/2 computer.

NetWare Server object

A leaf object that represents a server running NetWare on your network. A NetWare Server object can represent a server running any version of NetWare. Certain objects, such as Volume objects, refer to a NetWare Server object to help identify their locations.

See also Object.

NetWare user tools

Software that provides you with a graphical means of accessing network resources, such as volumes, directories, printers, and users. NetWare user tools allow you to perform tasks such as managing drive mappings, managing printer connections and setup, managing server connections, displaying network users, and sending messages. NetWare user tools are available for DOS, MS Windows, OS/2, Macintosh, and UNIX.

NetWare volume

A physical amount of hard disk storage space, fixed in size. A NetWare volume is the highest level in the NetWare directory structure (on the same level as a DOS root directory).

See also Volume.

NetWire

An online information service, which provides access to Novell product information, Novell services information, and time-sensitive technical information for NetWare users. NetWire is accessed through the CompuServe* Information Service. It requires a PC or compatible workstation, a modem, and a communications program.

Network

A group of computers that can communicate with each other, share peripherals (such as hard disks and printers), and access remote hosts or other networks. A NetWare network consists of workstations, peripherals, and one or more NetWare servers. NetWare network users can share the same files (both data and program files), send messages directly between workstations, and protect files with an extensive security system.

Network address

A network number that uniquely identifies a network cable segment; usually referred to as the *IPX external network number*.

See also IPX external network number.

Network backbone

A cabling system that NetWare servers and routers are attached to. The central cable handles all network traffic, decreasing packet transmission time and traffic on the network.

Network board

A circuit board installed in each workstation to allow stations to communicate with each other and with the NetWare server. Some printers contain their own network boards to allow them to attach directly to the network cabling. NetWare documentation uses the term *network board*. Documentation from other companies might use the terms *NIC* or *network card* instead of *network board*.

Network communication

Data transmission between workstations. Requests for services and data pass from one workstation to another through a communication medium such as cabling.

Network direct printer

Printers and third-party print queue servers that connect directly into the network.

Network drive

A common name for a *logical drive*.

See also Drive.

Network node

A personal computer or other device connected to a network by a network board and a LAN driver. A network node can be a server, workstation, router, printer, or fax machine.

Network number

A network number that uniquely identifies a network cable segment; usually referred to as the *IPX external network number*.

See also IPX external network number.

Network numbering

The system of numbers that identifies servers, network boards, and cable segments. These network numbers include the following:

- IPX external network number
- IPX internal network number
- Node number

See also IPS external network number; IPX internal network number; Node number.

Network printer

A printer shared in a network environment.

See also Printer.

Network supervisor

A generic term in NetWare 4 for the person responsible for configuring the NetWare server, workstations, user access (security), printing, etc. (Also referred to as *network administrator.*)

Network Support Encyclopedia Professional Volume

(NSE Pro) An electronic information database containing comprehensive information about network technology.

NETX

A VLM (NETX.VLM) in the NetWare DOS Requester that provides backwards compatibility with NETX and other older versions of the shell.

See also NetWare DOS Requester; Virtual Loadable Module.

NFS

Networked file system. NetWare® NFS* allows UNIX systems to integrate with NetWare 4 file systems.

See also NetWare Networked File System.

NIC

(Network interface card) A circuit board installed in each workstation to allow stations to communicate with each other and with the NetWare server. NetWare documentation uses the term *network board* instead of *NIC*.

NLM

(NetWare Loadable Module) A program you can load and unload from server memory while the server is running.

See also NetWare Loadable Module.

NLSP

(NetWare Link Services Protocol) A link-state routing protocol designed by Novell for IPX internetworks.

See also NetWare Link Services Protocol.

NNS

(NetWare Name Service) A naming service designed to provide more transparent access to resources in NetWare installations. NNS was a predecessor to NDS.

See also NetWare Directory Services; NetWare Name Service.

Node address

A number that uniquely identifies a network board; usually referred to as the *node number.*

See also Node number.

Node number

A number that uniquely identifies a network board, also known as *station address, physical node address,* and *node address.* Every node must have at least one network board, by which the node is connected to the network. Each network board must have a unique node number to distinguish it from all other network boards on that network.

See also IPX internetwork address.

Nonextended AppleTalk network

An AppleTalk network that does not support Phase 2 extensions, such as zone lists and network ranges.

See also AppleTalk Phase 2.

Normal (N) attribute

A file system attribute that indicates that no NetWare attributes are set.

See also Attributes.

NSE Pro

(NetWare Support Encyclopedia Professional Volume) A Novell database that contains comprehensive information about network technology.

Object

An NDS structure that stores information about a network resource (such as a user, group, printer, or volume). An object consists of categories of information, called *properties*, and the data in those properties. The information is stored in the NetWare Directory database.

Some objects represent physical entities, such as a user or a printer. Some objects represent logical entities, such as groups and print queues. Other objects, such as the Organizational Unit object, help you organize and manage objects.

An object is a structure where *information* about the entity is stored; it is not the actual entity.

Object rights

Qualities assigned to an object that control what the object can do with directories, files, or other objects.

See also Rights.

ODI

(Open Data-Link Interface) An architecture that allows multiple LAN drivers and protocols to coexist on network systems.

See also Open Data-Link Interface.

Open Data-Link Interface

(ODI) An architecture that allows multiple LAN drivers and protocols to coexist on network systems. ODI describes the set of interface and software modules used to decouple device drivers from protocol stacks and to enable multiple protocol stacks to share the network hardware and media transparently.

Optical disc

(Also *optical disk*) A form of removable media used to store data. An optical disc can be one- or two-sided. Some optical discs are read-only; others can be read from and written to. HCSS uses rewritable optical discs.

See also High Capacity Storage System.

Optical disc library

A high-capacity storage device, sometimes called a *jukebox*, that uses an autochanger mechanism to mount and dismount optical discs as needed.

See also High Capacity Storage System.

Organization object

A container object that helps you organize other objects in the Directory and allows you to set template information for users created in it. The Organization object is a level below the Country object (if used), and a level above the Organizational Unit object (if used).

See also Object.

Organizational Role object

A leaf object that defines a position or role within an organization. Use the Organizational Role object to specify a position that can be filled by different people, such as Team Leader or Vice President.

See also Object.

Organizational Unit object

A container object, a level below the Organization object, that helps you to further organize objects in the Directory and also allows you to set template information for users created in it.

See also Object.

OS/2 client

An OS/2 computer that connects to the network using NetWare Client for OS/2 software. The OS/2 client can store and retrieve data from the network and can run executable network files. OS/2 client workstations support IPX/SPX, NetBIOS, and Named Pipes to allow users access to OS/2-based applications such as the SQL Server.

See also NetWare Client for OS/2.

OS/2 Requester

A component of the software that connects OS/2 workstations to NetWare networks, allowing OS/2 users to share network resources.

See also NetWare Client for OS/2.

Packet

A unit of information used in network communication. Messages sent between network devices (such as workstations or NetWare servers) are formed into packets at the source device. The packets are reassembled, if necessary, into complete messages when they reach their destination. A packet might contain a request for service, information on how to handle the request, and the data to be serviced.

An individual packet consists of headers and a data portion. The different headers are appended to the data portion as the packet travels through the communication layers. A message that exceeds the maximum size is partitioned and carried as several packets. When the packet arrives at its destination, the headers are stripped off in reverse order and the request is serviced.

See also Communication protocols; Ethernet configuration; Large Internet Packet; NetWare Core Protocol.

Packet Burst protocol

A protocol built on top of IPX that speeds the transfer of multiple-packet NCP file reads and writes. The Packet Burst™ protocol speeds the transfer of NCP data between a workstation and a NetWare server by eliminating the need to sequence and acknowledge each packet.

Packet receive buffer

An area in the NetWare server's memory set aside to temporarily hold data packets arriving from the various workstations. The packets remain in this buffer until the server is ready to process them and send them to their destination. This ensures the smooth flow of data into the server, even during times of particularly heavy input/output operations.

See also Watchdog.

Paging

Allows NetWare 4 to assign memory noncontiguously. This is a feature of most high-powered CPU architectures (such as Intel* 80386/80486 and Pentium*).

Page tables are used to map physical addresses to logical memory. Each page table entry corresponds to a page in memory. A memory page is a 4KB block of RAM. A group of page tables is a *domain*.

Parallel port

A printer interface that allows data to be transmitted a byte at a time, all eight bits moving in parallel.

See also LPT ports.

Parent directory

The directory immediately above any subdirectory. For example, SYS:ACCTS would be the parent directory of the subdirectory SYS:ACCTS/RECEIVE.

See also File system.

Parent objects

Container objects that contain other objects.

See also Object.

Parent VLM

Acts as a multiplexor that routes calls to the correct child VLM. It ensures that each request to a child VLM reaches its destination.

See also Child VLM; Virtual Loadable Module.

Parity

A method of checking for errors in transmitted data.

See also Serial communication.

Partition (NetWare Directory)

A logical division of the NetWare Directory database.

See also NetWare Directory partition.

Partition (disk)

A logical unit into which NetWare server hard disks can be divided.

See also Disk partition.

Partition management

The method of managing NetWare Directory partitions and replicas. Partition management allows you to divide the Directory into partitions and manage various Directory replicas of these Directory partitions.

Password

The characters a user must type to log in. NetWare allows the network supervisor to specify whether passwords are required and, if so, to assign a login password to each user on the network. The network supervisor can also specify whether passwords must be unique and whether they must be changed periodically.

In NetWare 4, login passwords are encrypted at the workstation and put into a format that only the NetWare server can decode. This format helps prevent intruders from accessing network files.

See also Security; User object.

Path

The location of a file or directory in the file system. For example, the path for file REPORT.FIL in subdirectory ACCTG in directory CORP on volume SYS: of server ADMIN is

ADMIN\SYS:CORP\ACCTG\REPORT.FIL

See also Authentication; File system.

Physical memory

The RAM installed in a computer. NetWare servers use paging to address physical memory in 4KB blocks, or pages.

See also Logical memory; Memory protection; Paging.

Polled mode

A printer configuration option through which the port driver (NPRINTER) periodically checks, or polls, the data port to determine whether it is ready to accept data for transmission to the printer. The port's status is indicated by an electronic signal called a flag. Polling queries are made at each CPU timer tick (18 times per second).

Port (hardware)

A connecting component that allows a microprocessor to communicate with a compatible peripheral, such as a printer.

See also Parallel port; Serial port.

Port (software)

A memory address that identifies the physical circuit used to transfer information between a microprocessor and a peripheral.

Port driver

A driver that routes print jobs out of the print queue and through the proper port (for example, LPT1, LPT2, COM1) to the printer that will handle the job. The NPRINTER utility functions as a port driver in NetWare.

Postmaster

A user who has all of the following rights:

- Supervisor access to the NetWare MHS Messaging Server object
- Supervisor access to the Mailbox Location, Mailbox ID, and E-mail Address properties of users of the NetWare MHS messaging server
- Read access to the Message Routing Group that the NetWare MHS messaging server is in

Postmaster General

A user who has Supervisor access to the Message Routing Group that he or she resides in. A Postmaster General can add a messaging server to, or remove a messaging server from, the Message Routing Group.

You can assign several Postmasters General to a Message Routing Group.

Power conditioning

Methods of protecting sensitive network hardware components against power disturbances.

See also Uninterruptible power supply; UPS monitoring.

Primary time server

A server that synchronizes the time with at least one other Primary or Reference time server, and provides the time to Secondary time servers and to workstations.

See also Time synchronization.

Print device definition

A set of functions and modes found in a file with a .PDF extension that corresponds to a printer, plotter, or other peripheral. Print device definitions contain the necessary control sequences for setting or resetting the printer and for controlling bold, emphasis, italics, print size, font selection, colors, and other features, depending on the printer. A print device definition does not necessarily represent the full functionality of the printer. A print device definition can be modified to change the functions the machine can perform.

See also Print header and print tail.

Print driver

A driver that converts print jobs (usually generated by an application) to a format that can be read by the type of printer being used.

Print header and print tail

Contain transport control codes for the modes defined in PRINTDEF. The print header precedes the data to the print queue, and the print tail follows it. The default lengths are 64 and 16 bytes, respectively. These codes are especially critical for postscript printing.

Print job

A file, stored in a print queue directory, that is waiting to be printed. As soon as a print server sends a print job to the printer, the print job is deleted from the queue directory.

Print job configuration

A set of options that determine how a job is printed. Users can create print job configurations using the NetWare Administrator or PRINTCON.

Print queue

A network directory that stores print jobs. When the printer assigned to a print queue is ready, the print server takes the print job out of the print queue and sends it to the printer. The print queue can hold as many print jobs as disk space allows.

Print queue operator

A user who can edit other users' print jobs, delete print jobs from the print queue, or modify the print queue status by changing the operator flags. Print queue operators can also change the order in which print jobs are serviced. User ADMIN or equivalent can assign users to be print queue operators as necessary.

Print queue sampling interval

The time interval the print server waits between checking the print queues for jobs ready to be printed. The time period can be specified in PCONSOLE under "Printers." The default is 15 seconds.

Print server

A server that takes print jobs out of a print queue and sends them to a network printer.

Print Server object

A leaf object appearing in the Directory tree that represents a network print server.

See also Object.

Print Server operator

A user or member of a group delegated rights by User object ADMIN to manage the print server. A print server operator has rights to control notify lists, printers, and queue assignments.

Print tail

Contains transport control codes for the modes defined in PRINTDEF.

See also Print header and print tail.

Printer

A peripheral piece of hardware used to produce printed material.

See also Print device definition; Printing.

Printer form

A print option designed to prevent print jobs from being printed on the wrong paper. NetWare print services allows you to designate a printer form for a print job. The print jobs will not print if the correct paper is not in the printer.

Printer mode

A sequence of print functions (also called *printer commands*, *control sequences*, or *escape sequences*) that determines the appearance of the printed file. A printer mode can define the style, size, boldness, and orientation of the typeface. Print device modes are designated using NetWare Administrator or PRINTDEF.

Printer object

A leaf object that represents a physical printing device on the network.

See also Object.

Printing

The ability to transfer data from computer files to paper. NetWare 4 allows users to share printers on the network, where previously each personal computer had to have a printer attached to one of its printer ports. NetWare 4 uses a print queue and print server to allow workstations to print to a printer. The print server takes print jobs from the print queue and sends them to the printer.

See also Print queue; Print server; Printer.

Profile login script

A type of login script that sets environments for a group of users. Use profile login scripts if you have groups of users with identical login script needs. Profile login scripts are optional; if used, they execute after the container login script and before the user login script.

See also Login scripts.

Profile object

A leaf object that represents a login script used by a group of users who need to share common login script commands. The group of users do not need not to be located under the same container in the Directory tree. The group can also be a subset of users in the same container.

See also Object.

Prompt

A character or message that appears on the display screen and requires a response (such as a command or a utility name) from the user. Standard types of prompts include

- The DOS prompt, which, by default, displays the current drive letter followed by a > symbol (for example, F>)
- The OS/2 prompt, which, by default, displays the current drive mapping in brackets (for example, [C:\])
- The NetWare server console prompt, which displays a colon (:)

See also Drive mapping; Login scripts.

Property

A characteristic of an NDS object. Each type of object (such as a User object, Organization object, or Profile object) has certain properties that hold information about the object.

See also Object.

Property rights

Rights that apply to the properties of an NDS object.

See also Rights.

Protected mode

Provides the capability of multitasking (running more than one application or process at a time). Protected mode allocates memory to various processes running concurrently so that memory used by one process does not overlap memory used by another process.

See also Read/write replica.

Protocols

Conventions or rules used by a program or operating system to communicate between two or more endpoints.

See also AppleTalk protocols; Communication protocols; NetWare protocols and transports.

PUBLIC directory

The SYS:PUBLIC directory, created during network installation, that allows general access to the network and contains NetWare utilities and programs for network users.

See also File system; LOGIN directory; MAIL directory; SYSTEM directory.

Public files

Files that must be accessed by all NetWare users, including NetWare utilities, help files, and some message and data files. By convention, the files are located in SYS:PUBLIC for DOS users and in SYS:PUBLIC/OS2 for OS/2 users. All NetWare users have Read and File Scan rights to the files.

[Public] trustee

A special trustee that can be added to any object, directory, or file. By default, [Public] includes the Read right and the File Scan right. By making [Public] a trustee of an object, directory, or file, you effectively grant all objects in NDS rights to that object, directory, or file.

[Public] is only used in trustee assignments and must always be entered within square brackets. [Public] can be added or deleted like any other trustee. An Inherited Rights Filter blocks inherited rights for [Public] as it would any other trustee.

See also Inherited Rights Filter; Trustee.

Purge (P) attribute

A file system attribute that causes NetWare to purge the directory or file when it is deleted.

See also Attributes.

Queue

A network directory that stores each print job. When the printer assigned to the print queue is ready, the print server takes the print job out of the print queue and sends it to the printer.

See also Print queue.

Queue sampling interval

The time interval the print server waits between checking the print queues for jobs that are ready and waiting to be printed.

See also Print queue sampling interval.

Queue server mode

An operating mode used by many network-direct printers and hardware queue servers produced by various manufacturers. These devices either connect to a printer and then to the network or are installed in a port at the printer.

RAM

(Random Access Memory) The internal dynamic storage of a computer that can be addressed by the computer's operating system.

See also Memory.

Read-after-write verification

A means of assuring that data written to the hard disk matches the original data still in memory. If the data from the disk matches the data in memory, the data in memory is released. If the data does not match, the block location is recognized as bad, and Hot Fix redirects the data to a good block location within the Hot Fix Redirection Area.

See also Data protection.

Read-only replica

The Directory replica used to view, but not modify, Directory information.

See also NetWare Directory replica.

Read Only (Ro) attribute

A file system attribute that indicates that no one can write to the file.

See also Attributes.

Read right

A file system right that grants the right to open and read files. Also, a property right that grants the right to read the values of the property.

See also Rights.

Read/write replica

The Directory replica used to read or update Directory information (such as adding or deleting objects).

See also NetWare Directory replica.

Real mode

The mode that allows an 80286, 80386, or 80486 processor to emulate an 8086 processor and run as though it were an 8086 processor.

See also Protected mode.

Record locking

A feature of the NetWare operating system that prevents different users from gaining simultaneous access to the same record in a shared file, preventing overlapping disk writes and ensuring data integrity.

See also File locking.

Recursive copying

The process of copying a specified source directory to a destination directory until all files and subdirectories in and below the specified source directory are copied. Recursive copying copies all directories and files of a logical drive to the destination, keeping them exactly as they were in the source directory. The DOS and OS/2 XCOPY and BACKUP utilities use recursive copying, as does the NetWare NCOPY command.

Redirection area

(Hot Fix Redirection Area) The space on a hard disk set aside to hold data redirected from faulty disk blocks.

See Hot Fix.

Reference time server

A server that specifies the time to which all other time servers and workstations synchronize.

See also Time synchronization.

Remote boot

A method that allows a user to boot a workstation from remote boot image files on a NetWare server rather than from a boot diskette in the workstation's local drive. Client workstations that can start using remote booting do not need a floppy disk drive or a hard disk to function on the network; they are called *diskless workstations*. A diskless workstation relies on a Programmable Read-Only Memory (PROM) chip installed on its network board to communicate with the boot server.

See also Remote printer mode.

Remote connection

A connection between a LAN on one end and a workstation or network on the other, often using telephone lines and modems. A remote connection allows data to be sent and received across greater distances than those allowed by normal cabling.

Remote console

(RCONSOLE) Software that allows network supervisors to manage servers from a workstation. RCONSOLE enhances security since you can lock servers in a safe place and remove the keyboards and monitors.

Remote Program Load

(RPL) Technology based on the concept of storing an image of a bootable floppy diskette on a NetWare volume. Remote boot workstations use this image to start up the system prompt. These client workstations do not need a floppy disk drive or hard disk to function on the network; they are called *diskless workstations.* A diskless workstation relies on a Programmable Read-Only Memory (PROM) chip installed on its network board. This chip allows the workstation to communicate with the boot server.

Remote printer mode

An operating mode used by many network-direct printers and hardware queue servers produced by various manufacturers. Network-direct printers and hardware queue servers either connect to a printer and then to the network or are installed in a port at the printer.

Remote workstation

A terminal or personal computer connected to the LAN by a router or through a remote asynchronous connection. A remote workstation can be either a standalone computer or a workstation on another network.

Rename Inhibit (Ri) attribute

A file system attribute that prevents any user from renaming the directory or file.

See also Attributes.

Rename right

An object right that grants the right to change the name of an object, in effect changing the naming property.

See also Rights.

Replica

A copy of a NetWare Directory partition.

See also NetWare Directory replica.

Resource fork

The part of a Macintosh file that contains file resources, including Macintosh-specific information such as the windows and icons used with the file.

See also Data fork; Macintosh files.

Resource tags

Operating system tags that keep track of NetWare server resources such as screens and allocated memory. NLM programs request a resource from the NetWare server for each kind of resource they use and then give it a resource tag name. NLM programs return resources when they no longer need them. When the NLM is unloaded, the resources are returned to the NetWare server. Resource tags ensure that allocated resources are properly returned to the operating system upon termination of an NLM program.

Resources

The manageable components of a network, including

- Networking components, such as cabling, hubs, concentrators, adapters, and network boards
- Hardware components, such as servers, workstations, hard disks, and printers
- Major software components, such as the NetWare operating system and resulting network services (including file, mail, queue, and communication services)
- Minor software components that are controlled by the operating system of its subsystems—protocols, gateways, LAN and disk drivers, etc.
- Data structures and other network resources such as volumes, queues, users, processes, security, etc.

Restore

A retrieval of data previously copied and backed up to a storage media. Perform a restore if data has been lost or corrupted since the backup.

See also Backup; Data set.

Ribbon cable

A cable in which the wires are placed side by side in the insulation material instead of being bunched or twisted together in a circle inside the insulation material. Typically, ribbon cables are used for connecting internal disks or tape drives.

Rights

Qualities assigned to an object that control what the object can do with directories, files, or other objects. Creating, reading, and other operations can be done only if an object has rights to perform them.

Rights are granted to a specific directory, file, or object by *trustee assignments.* An object with a trustee assignment to a file, directory, or another object is a trustee of that file, directory, or object. Within each object is a list of who has rights to the object and what rights the object has to other objects. This list is the ACL property of the object. (Files and directories contain similar information, but not an ACL.)

See also Access Control List; Security.

RIP (IPX)

(Router Information Protocol) A protocol that provides a way for routers to exchange routing information on a NetWare internetwork.

See also Router Information Protocol.

Root directory

The highest directory level in a hierarchical directory structure. With NetWare, the root directory is the volume; all other directories are subdirectories of the volume.

See also Failure handling; File system.

[Root] object

An object in the Directory tree whose purpose is to provide a highest point to access different Country and Organization objects, and to allow trustee assignments granting rights to the entire Directory tree. Country, Organization, and Alias objects can be created at the [Root] object. The [Root] object is a place holder; it contains no information.

See also Directory tree; Object.

Router

A workstation or NetWare server running software that manages the exchange of information (in the form of data packets) between network cabling systems. A NetWare router runs as part of a NetWare server. It connects separate network cabling topologies or separate networks by way of the server's NetWare operating system.

Router Information Protocol

(RIP) A protocol that provides a way for routers to exchange routing information on a NetWare internetwork. RIP allows NetWare routers to create and maintain a database (or router table) of current internetwork routing information. Workstations can query the nearest router to find the fastest route to a distant network by broadcasting a RIP request packet.

Routers send periodic RIP broadcast packets containing current routing information to keep all routers on the internetwork synchronized. Routers also send RIP update broadcasts whenever they detect a change in the internetwork configuration.

See also Router; Service Advertising Protocol.

Salvageable files

Files saved by NetWare, after being deleted by users, that can be salvaged (recovered). Salvageable files are usually stored in the directory from which they were deleted. If the user deletes that directory, the file is saved in a DELETED.SAV directory located in the volume's root directory. Recovered files contain information about who deleted the files and when they were deleted.

Deleted files are saved until the user deliberately purges them or until the NetWare server runs out of disk allocation blocks on the volume. When the NetWare server runs out of blocks, it purges deleted files beginning with the files that were deleted first.

SAP

(Service Advertising Protocol) A protocol that provides a way for servers to advertise their services on a NetWare internetwork.

See also Service Advertising Protocol.

SBACKUP

A backup engine that provides backup and restore capabilities.

See also Storage Management Services.

SCSI

(Small Computer Systems Interface, commonly pronounced *scuzzy*) An industry standard that sets guidelines for connecting peripheral devices (such as hard drives and tape backup systems) and their controllers to a microprocessor. The SCSI interface defines both hardware and software standards for communication between a host computer and a peripheral device. Computers and peripheral devices designed to meet SCSI specifications have a large degree of compatibility.

SCSI bus

An interface that connects additional HBAs to controllers and hard disks.

See also Hard disk.

Search drive

A drive that the operating system searches when a requested file is not found in the current directory. Search drives are supported only from DOS workstations. A search drive allows a user working in one directory to access an application or data file located in another directory.

See also Drive mapping.

Search modes

Methods of operation that specify how a program uses search drives when looking for a data file.

Secondary time server

A server that obtains the time from a Single Reference, Primary, or Reference time server and provides the time to workstations.

See also Time synchronization.

Security

Elements that control access to the network or to specific information on the network. The six categories of security features are

- **Login security** controls which users can access the network.
- **Trustees** designate which users can access directories, files, or objects.
- **Rights** determine the level of access for each trustee.
- **Inheritance** passes rights from higher to lower levels.
- **Attributes** describe characteristics of directories and files.
- **Effective rights** list a user's actual rights to a directory, file, or object (including explicitly granted rights and inherited rights).

See also Attributes; Effective rights; Inherited Rights Filter; Rights; Security Equal To; Trustee.

Security Equal To

A property of every User object that lists other objects. The user is granted all rights that any object (such as User, Group, or Printer objects) in that list is granted, both to objects and to files and directories.

See also User object.

Seed router

A router that defines the range of network numbers for all routers in an AppleTalk network segment. Each AppleTalk network segment must have at least one seed router.

Semaphore

An integer value used to coordinate activities of both programs and processes to prevent data corruption in multiprocessor environments. Semaphores are used to synchronize interprocess communication by ensuring that certain event sequences do or do not occur.

See also File locking.

Serial communication

The transmission of data between devices over a single line, one bit at a time.

Serial port

A port that allows data to be transmitted asynchronously, one bit at a time. Typically, serial ports are used for modems or serial printers. On IBM PC-compatible computers, COM1 and COM2 are asynchronous serial ports.

Serialization

The process of serializing software to prevent unlawful software duplication.

Server

A computer in a network shared by multiple users.

NetWare server. A computer running the NetWare operating system software. (*See* NetWare server.)

Print server. A computer that takes print jobs out of a print queue and sends them to a network printer. (*See* Print server.)

Server console

The monitor and keyboard where you view and control NetWare server activity.

Server mirroring

An SFT III configuration which provides a secondary, identical server that can immediately take over network operations if the primary server fails.

Server protocol

Procedures that a NetWare server follows to accept and respond to workstation requests.

See NetWare Core Protocol.

Service Advertising Protocol

(SAP) A protocol that provides a way for servers to advertise their services on a NetWare internetwork. Servers advertise their services with SAP, allowing routers to create and maintain a database of current internetwork server information.

Routers send periodic SAP broadcasts to keep all routers on the internetwork synchronized. Routers also send SAP update broadcasts whenever they detect a change in the internetwork configuration.

See also Router; Router Information Protocol.

SFT

(System Fault Tolerance) A means of protecting data by providing procedures that allow you to recover from hardware failures. Three levels of SFT are available: Hot Fix, disk mirroring or duplexing, and server mirroring.

See also System Fault Tolerance.

Shareable (Sh) attribute

A file system attribute that allows a file to be accessed by more than one user at a time.

See also Attributes.

Short machine type

A four-letter (or less) name representing a brand of DOS workstations. The short machine type is similar to the long machine type; however, the short machine type is used specifically with overlay files.

See also Login scripts; Long machine type.

Single Reference time server

A server that provides time to Secondary time servers and to workstations. The Single Reference time server is the sole source of time on the network.

See also Time synchronization.

Small Computer Systems Interface

(SCSI) An industry standard that sets guidelines for connecting peripheral devices and their controllers to a microprocessor.

See also SCSI.

SMS

(Storage Management Services) A combination of related services that allow data to be stored and retrieved.

See also Storage Management Services.

SNA

(System Network Architecture) IBM's proprietary networking architecture; first introduced in 1974.

See also System Network Architecture.

Socket

The part of an IPX internetwork address, within a network node, that represents the destination of an IPX packet. Some sockets are reserved by Novell for specific applications. For example, IPX delivers all NCP request packets to socket 451h. Third-party developers can also reserve socket numbers for specific purposes by registering those numbers with Novell.

See also IPX internetwork address.

Source routing

A method used by IBM to route data across source-routing bridges. NetWare source-routing programs allow an IBM token ring network bridge to forward NetWare packets (or frames).

Source server

The server from which you migrate data files, bindery files, and other information to a NetWare 4 destination server during upgrade.

See also Destination server.

Sparse file

A file with at least one empty block. (NetWare will not write any block that is completely empty.) Databases often create sparse files.

SPX

(Sequenced Packet Exchange) A NetWare DOS Requester module that enhances the IPX protocol by supervising data sent out across the network. SPX verifies and acknowledges successful packet delivery to any network destination by requesting a verification from the destination that the data was received.

See also NetWare Client for OS/2; NetWare DOS Requester; NetWare user tools.

STARTUP.NCF

A NetWare server boot file that loads the NetWare server's disk driver and name spaces and some SET parameters.

See also Boot files.

Station

Usually a shortened form of *workstation*, but can also be a server, router, printer, fax machine, or any computer device connected to a network by a network board and a communication medium.

Station address

A number that uniquely identifies a network board; usually referred to as the *node number.*

See Node number.

Stop bit

A signal that indicates the end of a character.

See also Serial communication.

Storage device

A device used to back up data from a server or workstation. An example of a storage device is an external tape drive that backs up data from a hard disk to magnetic tape.

Storage Management Services

(SMS) Services that allow data to be backed up and restored. SMS is independent of backup/restore hardware and file systems (such as DOS, OS/2, Macintosh, MS Windows, or UNIX).

See also Backup hosts and target; NetWare Directory database; NetWare Loadable Module; Storage Management Services; Target Service Agent.

STREAMS

An NLM program that provides a common interface between NetWare and transport protocols (such as IPX/SPX, TCP/IP, SNA, and OSI) that need to deliver data and requests to NetWare for processing. By making the transport protocol transparent to the network operating system, STREAMS™ allows services to be provided across the network, regardless of the transport protocols used.

Subdirectory

A directory below another in the file system structure. For example, in SYS:ACCTS\RECEIVE, RECEIVE is a subdirectory of SYS:ACCTS.

See also File system.

Subnetwork mask

Indicates how the host portion of the IP address is divided into subnetwork addresses and local host address portions. The network mask is a 32-bit number with all ones for all network and subnetwork address portions of the complete IP address, and all zeros for host address portions.

With a 16-bit Class B IP network address, a 4-bit subnetwork address, and a 12-bit host address, the subnetwork mask consists of 20 ones and 12 zeros. In essence, a subnetwork mask locally extends the network address portion of an IP address.

Subordinate replica

A Directory replica that is automatically placed on a server if the parent Directory partition has a master, read/write, or read-only replica and the child Directory partition does not. Subordinate replicas cannot be modified.

Supervisor right

A file system right that grants all rights to the respective directory and files. Also, an object right that grants all access privileges to all objects. Also, a property right that grants all rights to all properties or to selected properties.

See also Rights.

Synchronization

Replica synchronization. A means of ensuring that replicas of a Directory partition contain the same information as other replicas of that partition. (*See* NetWare Directory replica.)

Time synchronization. A method of ensuring that all servers in a Directory tree report the same time. (*See* Time synchronization.)

System (Sy) attribute

A file system attribute that marks directories or files for use only by the operating system.

See also Attributes.

SYSTEM directory

The SYS:SYSTEM directory, created during network installation, that contains NetWare operating system files as well as NLM programs and NetWare utilities for managing the network. By default, user ADMIN (or a user with ADMIN equivalent rights) has rights to the SYS:SYSTEM directory. Do not delete the SYSTEM directory.

See also File system; LOGIN directory; MAIL directory; PUBLIC directory.

System Fault Tolerance

(SFT) A means of protecting data by providing procedures that allow you to automatically recover from hardware failures.

See also Data protection.

System login script

In NetWare 2 and NetWare 3, a type of login script that sets general environments for all users. In NetWare 4, the container login script replaces the system login script.

See Container login script; Login scripts.

System Network Architecture

(SNA) IBM's proprietary networking architecture; first introduced in 1974. SNA is the technology that allows LAN systems to be connected to IBM mainframe computers.

Tape backup device

An internal or external tape drive that backs up data from hard disks.

Target

Any server, workstation, or service on the network that has a Target Service Agent loaded. A target can have its data backed up or restored. If you are backing up and restoring to the host server, the target and the host are the same.

See also Host; Target Service Agent.

Target Service Agent

A program that processes data moving between a specific target and an SMS-compliant backup engine, such as SBACKUP.

See also Backup hosts and target; Storage Management Services.

TCP/IP

(Transmission Control Protocol/Internet Protocol) An industry-standard suite of networking protocols that enables dissimilar nodes in a heterogenous environment to communicate with one another. TCP/IP is built upon four layers that roughly correspond to the seven-layer OSI Reference Model.

Time synchronization

A method of ensuring that all servers in a Directory tree report the same time. Clocks in computers can deviate slightly, resulting in different times on different servers. Time synchronization corrects these deviations so that all servers in a Directory tree report the same time and provide a *time stamp* to order NDS events.

Topology

The physical layout of network components (such as cables, stations, gateways, and hubs). Three basic topologies are used:

- **Star network**. Workstations are connected directly to a NetWare server but not to each other.
- **Ring network**. The NetWare server and workstations are cabled in a ring; a workstation's messages may have to pass through several other workstations before reaching the NetWare server.
- **Bus network**. All workstations and the NetWare server are connected to a central cable (called a *trunk* or *bus*).

Transaction Tracking System

(TTS) A system that protects data from corruption by backing out incomplete transactions which result from a failure in a network component. When a transaction is backed out, data is returned to the state it was in before the transaction began. TTS is a standard feature on NetWare 4 servers.

See also Cache memory.

Transactional (T) attribute

A file system attribute that indicates the file is protected by TTS.

See also Attributes.

Transmission Control Protocol

(TCP) An industry-standard suite of networking protocols that enables dissimilar nodes in a heterogenous environment to communicate with one another.

See also TCP/IP.

Trustee

A user or group granted rights to work with a directory, file, or object; the object is called a *trustee* of that directory, file, or object. Rights are granted to objects (making them trustees) by *trustee assignments.* Trustee assignments are part of the directory, file, or object to which they grant access. Trustee assignments are stored in a *trustee list.* An object's trustee list is stored in the object's ACL property.

See also Inherited Rights Filter; Security.

TTS

(Transaction Tracking System) A system that protects database applications from corruption by backing out incomplete transactions that result from a failure in a network component.

See also Transaction Tracking System.

Turbo FAT index

A special FAT index used when a file exceeds 64 blocks (and the corresponding number of FAT entries). NetWare creates a turbo FAT index to group together all FAT entries for that file. The first entry in a turbo FAT index consists of the first FAT number of the file. The second entry consists of the second FAT number, and so on. A turbo FAT index enables a large file to be accessed quickly.

See also File Allocation Table.

Unbinding

The process of removing a communication protocol from network boards and LAN drivers.

See also Binding and unbinding.

Uninterruptible power supply

(UPS) A backup power unit that supplies uninterrupted power if a commercial power outage occurs.

See also UPS monitoring.

UNIX client

A UNIX computer connected to the network. The UNIX client stores and retrieves data from the NetWare server and runs executable network files. The UNIX client provides multiple NetWare-client multitasking on a single station. UNIX clients include IPX/SPX and NCP/IPX communication protocols to allow other NetWare clients access to UNIX applications.

See also Client.

Unknown object

A leaf object that represents an NDS object that has been corrupted and cannot be identified as belonging to any of the other object classes. After migrating to NetWare 4 from NetWare 2 or 3, bindery objects might appear as Unknown objects.

See also Object.

Unloading

The process of unlinking NLM programs from the NetWare operating system.

See also Loading and unloading.

Upgrade

The process of converting your network to NetWare 4 from any earlier version of NetWare or from another network operating system.

UPS

(Uninterruptible power supply) A backup power unit that supplies uninterrupted power if a commercial power outage occurs.

See also Uninterruptible power supply.

UPS monitoring

The process a NetWare server uses to ensure that an attached UPS is functioning properly. A Novell-certified UPS is attached to a server to provide backup power. (You can also attach a UPS to workstations without installing UPS monitoring hardware on the stations.) When a power failure occurs, NetWare notifies users. After a timeout specified in SERVER.CFG, the server logs out remaining users, closes open files, and shuts itself down.

See also Uninterruptible power supply.

User login script

A type of login script that sets environments specific to a user. Use user login scripts to contain items that cannot be included in system or profile login scripts. User login scripts are optional; if used, they execute after system and profile login scripts.

See also Login scripts.

User object

A leaf object in NDS that represents a person with access to the network. A User object stores information about the person it represents.

See also Accounting; Group object; Login scripts; Security Equal To.

User object ADMIN

A User object, created automatically during NetWare 4.1 installation, that has rights to create and manage objects. When you first create the Directory tree, ADMIN is given a trustee assignment to the [Root] object. This trustee assignment includes the Supervisor object right, which means that ADMIN has rights to create and manage all objects in the tree.

See also User object.

User template

A file containing default information you can apply to new User objects to give them default property values. This helps if you are creating many users who need the same property values.

Utilities

Programs that add functionality to the NetWare operating system. NetWare 4 utilities support MS Windows, OS/2, and DOS environments. (See *Utilities Reference.*)

See also NetWare Loadable Module; NetWare user tools.

VDT

A table that keeps track of volume segment information.

See also Volume Definition Table.

Virtual Loadable Module

(VLM) A modular executable program that runs at each DOS workstation and enables communication with the NetWare server.

See also NetWare DOS Requester.

VLM

A modular executable program that runs at each DOS workstation and enables communication with the NetWare server.

See also Virtual Loadable Module.

Volume

A physical amount of hard disk storage space, fixed in size. A NetWare volume is the highest level in the NetWare file system (on the same level as a DOS root directory). In NDS, each volume is also a Volume object in the Directory.

See also File system; Volume Definition Table.

Volume Definition Table

(VDT) A table that keeps track of volume segment information such as volume name, volume size, and volume segments locations on various network hard disks. Each NetWare volume contains a VDT in its NetWare partition.

See also Volume.

Volume object

A leaf object that represents a physical volume on the network.

See also Object; Volume.

Volume segments

A physical division of a volume. Different segments of a volume can be stored on one or more hard disks, allowing you to create large volumes. A single hard disk can contain up to eight volume segments belonging to one or more volumes, and each volume can span up to 32 segments.

See also Volume.

Wait state

A period of time when the processor does nothing; it simply waits. A wait state is used to synchronize circuitry or devices operating at different speeds.

Wait time

In a NetWare UPS system, the number of seconds the UPS waits before signaling to the NetWare server that the normal power supply is off. The NetWare server then alerts attached workstations to log out.

See also Uninterruptible power supply.

WAN

(Wide area network) A network that communicates over a long distance, such as across a city or around the world.

See also Wide area network.

Watchdog

Packets used to make sure workstations are still connected to the NetWare server. All settings are determined by the SET parameters.

See also Packet receive buffer.

Wide area network

(WAN) A network that communicates over a long distance, such as across a city or around the world. A local area network becomes a part of a wide area network when a link is established (using modems, remote routers, phone lines, satellites, or a microwave connection) to a mainframe system, a public data network, or another local area network.

See also Local area network.

Workstation

A personal computer connected to a NetWare network and used to perform tasks through application programs or utilities. Also referred to as a *client* or shortened to *station*.

See also Client.

Write right

A file system right that grants the right to open and write to files. Also, a property right that grants the right to add, change, or remove any values of the property.

See also Rights.

Zones

Arbitrary groups of nodes on an AppleTalk internetwork. Zones provide divisions in a large internetwork. Each node belongs to only one zone at a time. The zone that a node belongs to is determined automatically when that node connects to the network.

See also AppleTalk protocols.

Notes